Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

Firewalls, Perimeter Protection and Virtual Private Networks

Practical Assignment

New England SANS
5 – 9 September 2001

Steve Greenham

1   Assignment 1 – Security Architecture

Define a security architecture for GIAC Enterprises, an e-business which deals in the online sale of fortune cookie sayings. Your architecture must include the following components:

- filtering routers
- firewalls
- VPNs to business partners
- secure remote access; and
- internal firewalls.

Your architecture must consider access requirements (and restrictions) for:

- Customers (the companies that purchase bulk online fortunes)
- Suppliers (the authors of fortune cookie sayings that connect to supply fortunes)
- Partners (the international partners that translate and resell fortunes).

Include a diagram or set of diagrams that shows the layout of GIAC Enterprises' network and the location of each component listed above. Provide the specific brand and version of each perimeter defense component used in your design. Finally, include an explanation that describes the purpose of each component, the security function or role it carries out, and how the placement of each component on the network allows it to fulfill this role.

1.1 Business Requirements

We are not given any information about the size of GIAC Enterprises; how many staff it employs, its location(s), its turnover or the volume of transactions it processes each day. Similarly we are not given any information about the number of customers, suppliers or partners with whom GIAC Enterprises interacts. Before we can start designing a security architecture we must have a much better understanding of the business it is to support. Consequently our first step is to research the company; firstly by studying its annual report and then by visiting its offices and interviewing key business stakeholders.

1.1.1 GIAC Enterprises Company Synopsis

GIAC Enterprises was founded in 1999 and is a private limited company based in Ware, Hertfordshire, UK. From small beginnings above a Chinese restaurant on Ware High Street it has grown to employ a staff of 30 people and last year moved into a suite of offices in a business park on the edge of town. Responsibility for information systems lies with the accounts department who sent their most IT-literate person on Microsoft NT workstation, NT server, IIS and SQL server courses. She would have completed her MCSE in NT4.0 except that Microsoft retired the qualification. GIAC Enterprises management privately admitted they are relieved about this because they were expecting her to leave to get a full-time job in IT once she had her MCSE.

The company initially maintained its collection of fortune cookie sayings in a Microsoft Access database that was created by the owner's son as part of his GCSE Information

Technology project and initially populated with 5,000 fortune cookie sayings. The database has grown by around 100 sayings a day and currently holds around 70,000 sayings, now in a SQL server database accessed by up to 20 simultaneous users. The IT expert in accounts performed the migration from a standalone Access database to LAN and SQL server. Internet access is currently achieved by ISDN dial-up to a local ISP from the NT server which provides proxy services and POP3 mailboxes using Mailgate v3.5[i] GIAC Enterprises also have Axent Intruder Alert running on some of their servers, although nobody routinely checks the logs.

At present sayings are received by email and cut/pasted into the database. Customers bulk-purchase sayings that are manually extracted from the database and emailed as text files. Recently a number of sayings have turned up in competitors' fortune cookies and the owners are concerned that email is being intercepted. Although they have experimented with PGP it has caused a large administrative and support overhead. The majority of customers are unwilling to install PGP just to deal with GIAC Enterprises.

1.1.2 Proposed Business Model

GIAC Enterprises want to become an e-Business. They see an opportunity in making their database available on the Internet to avoid the time-consuming process of manually processing emails. The present assignment is to design a suitable system to support the proposed business model.
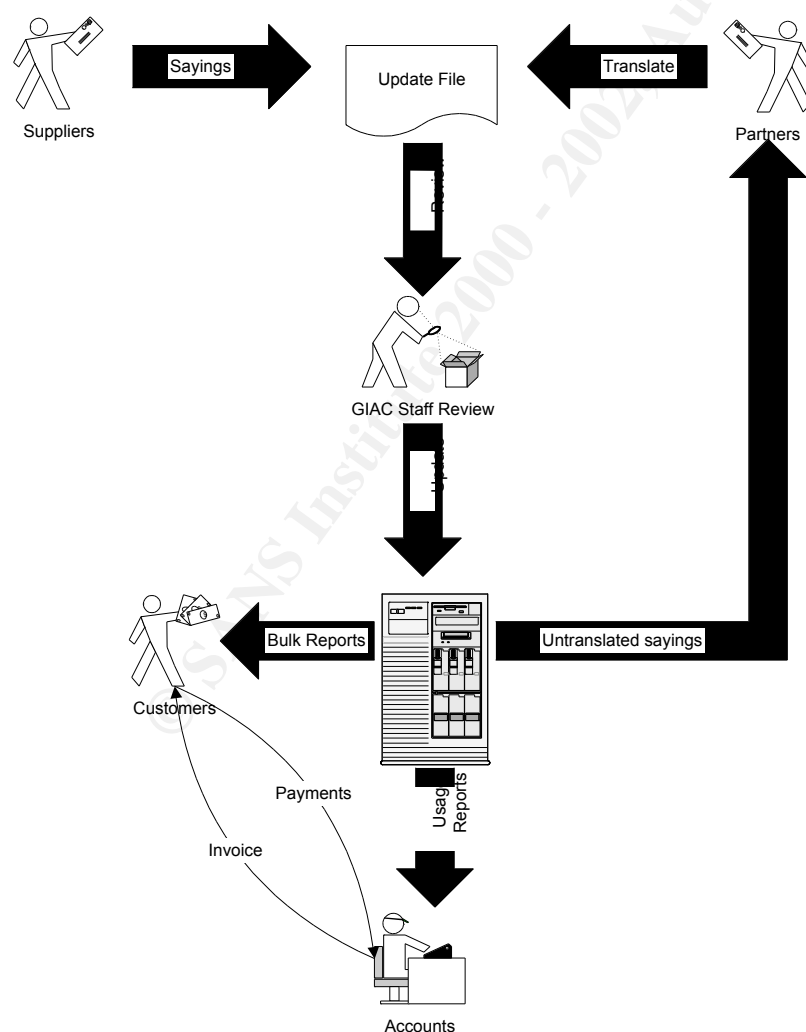
Interaction between stakeholders will be as follows:-

1.1.2.1 Customers

Customers will access a secure web site, authenticating with a two-factor mechanism such as SecurID. They will be able to browse fortune cookie sayings by category, create customised lists and download them as text files. They will be charged for access to the system and for each cookie saying that they view or download. This gives them the right to use the saying in their cookies.

The database needs to be able to track which customers have seen which sayings, allowing them free access to sayings they have already seen and paid for, while billing them for new sayings. Each customer will have a credit limit and will only be able to access new sayings while they remain within their credit limit. This offers some protection against customers downloading the entire database.

VPN solutions are required between customers and GIAC Enterprises because of the risk that bulk transfers of sayings may be intercepted.

1.1.2.2 Suppliers

Suppliers are the authors of fortune cookie sayings. GIAC Enterprises has a very open policy to authors and anyone who wishes can apply to become an accredited fortune cookie saying provider.

Suppliers log in to the web site to enter their sayings. This will be achieved through a cgi script that displays a form into which the supplier types the saying. The form appends the saying to a text file on the web server that is periodically used to update the database.

Suppliers will only be able to view the sayings that they entered and will not have a facility to bulk view or download sayings. They will be paid a fixed fee for each saying they submit once it is accepted for entry to the database.

1.1.2.3 Partners

Translation of fortune cookie sayings into other languages for distribution through partner organisations is a new opportunity. It has not been feasible through the existing email-based system because it requires partners to have online access to the database so that they can view, translate and store the fortune cookie sayings.

Partners will authenticate with the web server through user name and password. The database will track the language(s) in which each partner is fluent and will only display sayings that have not yet been translated into those languages. Sayings will be displayed one at a time and a translation must be entered in order for the next to be displayed.

The translations appended to the batch update file as is the case for suppliers. Again, GIAC staff will validate them before they are added to the database and payment authorised.

Some partners are simply translators while others are foreign fortune cookie companies that both translate and buy sayings. Bulk purchase of foreign language sayings will be handled through the customer interface described above.

1.1.2.4   GIAC Enterprises Staff

GIAC Enterprises back office staff will be responsible for viewing fortune cookie sayings and approving them for addition to the database. The update files created by suppliers and partners will be added to a temporary table on the database until a GIAC Enterprises member of staff has approved them. Approval causes the saying to be transferred to the live table and the supplier to be credited. As with partners, the workflow capabilities of the database need to display sayings in the languages in which the members of staff are proficient.

GIAC Enterprises accounts staff will be responsible for making payments to suppliers and partners, collecting payments from customers and partners and tracking usage of the system in order to set appropriate credit limits and report to management. They will also be responsible for generic financial functions such as payroll.

GIAC Enterprises intends to offer its resident IT expert the role of IT Manager with responsibility for the ongoing maintenance of the system. She will work with a contract IT Solutions provider to implement the system we design.

1.2   Security Architecture

1.2.1   Design Principles

1.2.1.1   SANS Top Twenty Internet Risks

In considering our security architecture we have borne in mind the business requirements of GIAC Enterprises, already discussed, and the SANS/NIPC consensus document of the Top 20 Internet security risks[ii]. In brief, the relevant generic and Windows vulnerabilities are

- **Default Installs of Operating Systems and Applications**, leaving services running that are not needed and can be exploited. We address this by hardening systems acting as peripheral components so that they only run the required services.
- **Missing/Weak passwords.** We address this by using two-factor authentication for external/privileged use and a strong password policy for internal use.
- **Backups**: Although not specifically addressed in this document, all systems will be backed up daily.
- **Large number of open ports.** Only those ports that are required will be opened. Ports not in use will be stealthed.
- **Not filtering packets for correct incoming and outgoing addresses.** This is addressed at length in the border router configuration.
- **Non-existent or incomplete logging**. Peripheral component logs are consolidated to a syslog server from where activity reporting is performed. The

approach to reporting is to filter out expected activity and then manually review what remains.

- **Vulnerable CGI Programs.** Only those programs necessary for operation of the GIAC web server are installed. CGI programs, scripts and example code that is not needed are not installed.
- **Unicode Vulnerability (Web Server Folder Traversal).** The IIS5.0 web server is patched to Service Pack 2 to remove this vulnerability.
- **ISAPI Extension Buffer Overflows.** The IIS5.0 web server is patched to Service Pack 2 to remove this vulnerability.
- **NETBIOS - unprotected Windows networking shares** and **Information leakage via null session connections** NetBIOS ports blocked at the external router.
- **Weak hashing in SAM (LM hash)** NTLMv2 is used throughout the GIAC internal network. Password audits are run quarterly using LC3[iii] and users informed of weak passwords.

1.2.1.2 Resource Constraints and Scaleability

GIAC Enterprises has undergone rapid growth over the past year and, although still a small company, continues to grow rapidly. For this reason we have factored in capacity for four-fold increases in hosts and transactions without significantly changing the network architecture.

The Company has agreed to allocate a full time headcount to IT support but the incumbent is relatively inexperienced in network administration. Therefore we have partitioned the network functions into distinct private address ranges and used simple class C masks. For the same reason, we have standardised on Windows 2000 platforms wherever possible. This reduces the risk of an unfamiliar *nix based system being misconfigured.

1.2.1.3 Physical Security

Perimeter components, management console, external, third party, service and secure networks are all physically located in the machine room. Only the internal subnet will extend beyond the machine room. As the other subnets each have few hosts, consideration could be given to housing them in separate cabinets for each subnet.

Physical access to the machine room will be restricted to authorised staff with entry/exit to the room logged.

1.2.1.4 Managing Change

The closer to the perimeter we get the less change will be allowed and the more tightly it will be controlled; so configuration changes to the border router should be very rare and made only after consultation and consideration. Firewall rule changes may be made more frequently but formal change control will still be followed and, other than in emergencies, changes only made after testing outside business hours. User-level changes, such as to give staff access to the Internet, may be made at any time by adding them to the appropriate Active Directory access group.

1.2.2 Network Topology

Our architecture provides five distinct subnets; an external network which links the firewall to the border router, a third party network on which VPN connections terminate,

a service network which hosts all publicly accessible servers, an internal company
network and a secure network on which particularly sensitive servers are hosted.

Address	Mask	Function
192.168.40.0	255.255.255.0	External ("red") network linking external firewall interface to border router.
192.168.10.0	255.255.255.0	Third Party Network. All third party VPN connections will terminate on this network segment.
192.168.20.0	255.255.255.0	Service ("yellow") Network. This network segment will house all hosts that are accessible from outside the company, whether by VPN connections or directly from the Internet.
192.168.30.0	255.255.255.0	Internal ("green") Network: All internal clients and non-sensitive servers will be on this network. The following network address convention is recommended :- 192.168.30.1 – 192.168.30.127 Windows 2000 clients. Addresses dynamically allocated by DHCP server (running on Windows 2000 domain controller) 192.168.30.193 – 192.168.30.207 Windows 2000 servers with static IP addresses. 192.168.30.222 Windows 2000 management console with static IP address
192.168.50.0	255.255.255.0	Secure Network: Hosts on this network will be those for which we require additional protection, possibly because they hold sensitive data (e.g. financial and employee records) or because they have a security function (e.g. syslog server)

Each of these networks has a distinct level of trust that will be reflected in our firewall
rules.

o The external network is least trusted. It can only initiate connections to the service
  network.
o The service network is second least trusted, can only initiate connections to the
  external network, with the exception of web server being allowed to request data
  from data server on the secure network.
o The third party subnet is next most trusted; it can initiate connections to service
  network and to the internal network (allowing remote staff to log in) We may wish
  to restrict access to the secure network from VPN connections.
o The internal subnet is next most trusted. It can access all subnets, but setting rules
  for specified clients, servers and protocols more tightly restricts that facility.
o The secure subnet is the most tightly controlled and does not have any interactive
  clients. Consequently traffic through the internal firewall can be tightly restricted to
  specific activities and protocols.

There will be no default route from the internal network to the Internet. Instead, all Internet access must be made via a proxy server (MS Proxy Server 2 running on Windows 2000).

Perimeter components will be configured so that all traffic will be denied except that which is specifically allowed.

1.2.3 External Router

We have adopted the principle of defence in depth. The outer layer, the border router, will be configured as a simple packet filter whose function is primarily to limit the amount of inbound traffic reaching the firewall and to prevent outbound address spoofing. We have selected a Cisco 1720 Modular access router[iv] to provide this functionality because this model provides a reasonable trade-off between price and flexibility. WAN interface cards are available for a range of Internet connections; from low speed serial and ISDN, through ADSL to T1.

1.2.4 Main Firewall

We have selected the Symantec Enterprise Firewall v6.5 (formally Raptor) because it integrates well with the host based intrusion detection system (Axent IntruderAlert) that was already in place. It will run on a hardened Windows 2000 Advanced Server with four network interface cards serving the External, Third Party, Service and Internal networks respectively. Ideally all NICs would be 10mbit to protect the firewall server from excess traffic, however it can be difficult to source these now that 10/100 cards are more common.

Network Address Translation will be used to convert between the RFC1918 private IP address ranges used on the GIAC Enterprises private subnets and the public IP addresses allocated to GIAC by their ISP.

1.2.5 VPN Security Gateway

Three options were considered for VPN access; the first was to route the traffic through the border router and Raptor firewall. Although this has the benefit of using existing components, so reducing cost, it was rejected because it introduced complexity to the firewall configuration rules and also significantly increased load on the firewall.

The second option was to use the same Internet link but have the border router direct VPN traffic to a separate security gateway on a third party subnet. This avoided the cost of a dedicated Internet link for third party traffic while retaining the benefit of routing third party traffic in to a separate firewall interface. After some consideration this approach was rejected because the Interface link and router then become single points of failure for both generic internet access and third party access.

Our recommended option is to provide a separate Internet link from the Internet to a Raptor PowerVPN security gateway on the third party network.Terminating on a different subnet keeps third party traffic separate making management, logging and monitoring easier. Note that traffic on the third party subnet will be in the clear, it is decrypted at the security gateway. The VPN gives us security while data is being transmitted over the public Internet, not once it arrives on the GIAC network.

Having data arrive at the main firewall on a separate network interface makes it easier to configure policies to allow access to the Internal network and service network. Similarly risks introduced by compromised remote clients allowing third party Internet users to tunnel through to the GIAC internal network are mitigated because the tunnel is only to the third party network, still outside the firewall.

The separate gateway, with its own address, allows us to route incoming traffic completely independently of the traffic to the external network. This provides an opportunity for redundancy with the traffic being carried by a different ISP. This introduces a lot of flexibility and an upgrade path as traffic increases and the links are upgraded so that eventually both normal Internet and VPN traffic could be routed over diverse high-speed links.

1.2.5.1   Remote Access By Staff

Staff using the VPN gateway to connect to GIAC Enterprises' internal network will use the Raptor remote client RaptorMobile 6.5 on their Windows 2000 laptop PCs and connect through the same Internet Service Provider as the VPN gateway connection. Use of the same ISP reduces latency by, hopefully, minimising the number of hops although service level cannot be guaranteed when using the public Internet.

Our security policy does not allow use of home PCs for remote access because of the risk of virus infection and GIAC intellectual property being stored on uncontrolled clients. Configuration management and audit of GIAC laptops, including update of antivirus software, will be with Microsoft Systems Management Server each time the laptop connects.

Staff will use SecurID tokens when accessing GIAC Enterprises remotely, this provides two factor authentication – something they have (the SecurID card, which displays a different number each minute) and something they know (an assigned PIN number). Systems admin staff will also use SecurID authentication when administering service network hosts from the management console. This provides an additional layer of security, and if we've paid to install an ACE server we should get the maximum leverage from it.

1.2.5.2   Remote Access By Third Parties

The most significant third party constituency are the customer organisations who download fortune cookie sayings in bulk. We will allow them to connect directly to our service network through the VPN gateway using network to network tunnels. Unlike our own staff, we have less control over the VPN software or ISP they use so we have to rely on IPSEC standards to ensure interoperability. This is discussed in depth later.

Third party connections will only be allowed to access the service network so we will not insist on two factor authentication at this time – but we may revise this policy for our major customers and issue them with SecurID cards or software in the future.

1.2.6 Internal Firewall

An internal firewall separates our secure network from the internal network. This allows access to more sensitive hosts to be more tightly controlled and monitored than if all internal hosts were on the same network. A dedicated firewall appliance will be used at

this point – a Nokia IP 330 running Checkpoint Firewall-1 v4.1 SP4 (build 41864). Use of a different brand of firewall ensures that any vulnerability in the Raptor firewall cannot also be used against the internal firewall.

1.2.7 Intrusion Detection

As we said in the introduction, GIAC Enterprises already have Symantec (Axent) Intruder Alert installed. This is a host-based intrusion detection system in which all sensitive hosts have agents installed reporting back to a centralised management console. The Intruder Alert rule sets installed add to Windows' generic event logging by allowing rule templates to be applied for particular operating systems and services. Intruder Alert can detect system events such as users or administrators logging in, critical file checksums changing, port scans etc. When such rules are triggered it can react by making an entry in its log, firing off SNMP traps, sending email or pager alerts. Intruder Alert includes a rule-set for Raptor firewalls and this would be applied to our firewall host.

GIAC Enterprises did not have a network based intrusion detection system so, subject to financial constraints, we recommend Symantec Netprowler v3.5. This gives a single vendor for host and network intrusion detection as well as VPN and Firewall ensuring maximum interoperability.

1.2.8 Other Hosts and Services

The following hosts and services form part of the GIAC Enterprises network and provide a supportive role to the components primarily responsible for perimeter defense.

1.2.8.1 Time Server

A GPS-linked network time protocol server is on the service network and all GIAC systems are configured to synchronise their clocks with it on a regular basis. This ensures that all logs are synchronised and events can be cross-referenced between different systems' logs.

1.2.8.2 SMTP Server

The SMTP Server is located on the services network from where it can accept and send email via the firewall's external interface. It hosts POP3 mailboxes for GIAC staff so incoming mail is held here for a time before it is collected by individuals' mail clients.

1.2.8.3 DNS Service

The Domain Name Service is provided by the Raptor firewall. This is because only a limited number of hosts, less than 100, are currently in use and the majority of these are clients with DHCP allocated IP addresses. Should GIAC Enterprises grow to a size where hosting the DNS on the firewall server is unsustainable a split DNS could be implemented. In such a case the external DNS server would be located on the service network and internal DNS server on the internal or secure networks.

1.2.8.4 Web Server

The GIAC Web server is Microsoft Internet Information Server 5.0 running on Windows 2000. It provides the publicly accessible web site (http://www.giac.co.uk) as well as the secured web site only accessible by customers and GIAC staff

(https://fortunes.giac.co.uk). The secured site is configured such that it only accepts HTTPS requests from the third party and internal networks.

1.2.8.5 Proxy Server

A proxy server on the internal network handles HTTP and FTP from internal clients. This conserves network bandwidth by caching popular pages locally, provides a point of authentication for staff accessing the Internet and maintains comprehensive logs of Internet use (and abuse). It is also a point at which blocking software could be installed to prevent access to inappropriate sites.

1.2.8.6 Syslog Server

A server on the secure network is used to consolidate the logs of all peripheral components allowing activity to be tracked and reported across the different systems.

1.2.8.7 SecurID ACE Server

The SecurID server is located on the secure network and responsible for authenticating those clients using SecurID tokens. Essentially the server is passed the two numbers entered by the client (number displayed on the token and PIN) and confirms whether these are correct. The number displayed on the token changes each minute.

1.2.9 Network Diagram

Internet

Frame Relay

512 kbps

Partner

Cisco 1720 Border Router

192.168.40.254

External Network 192.168.40.0/24

GPS

NTP Server
192.168.20.202

192.168.40.10

Service Network 192.168.20.0/24

192.168.10.10   192.168.20.10

Third Party Network 192.168.10.0/24

192.168.30.10

WWW
Server
192.168.20.201

External
DNS
Server
192.168.20.253

SMTP
Server
192.168.20.200

Internal Network 192.168.30.0/24

Windows 2000 workstations
x30

DHCP 192.168.30.128 -192

Management
Console
192.168.30.222

F & P
Server
192.168.30.207

Domain
Controller
192.168.30.208

Proxy
Server
192.168.30.209

Internal
DNS
Server
192.168.30.253

Secure Network
192.168.50.0/24

ACE Server
192.168.50.195

Database
Server
192.168.50.200

Syslog
Server
192.168.50.210

2   Assignment 2 – Security Policy

Part 1 – Define Your Security Policy

Based on the security architecture that you defined in Assignment 1, provide a security
policy for AT LEAST the following three components:

- Border Router
- Primary Firewall
- VPN

You may also wish to include one or more internal firewalls used to implement defence
in depth or to separate business functions.

By 'security policy' we mean the specific Access Control List (ACL), firewall ruleset,
IPSec policy, etc. (as appropriate) for the specific component used in your architecture.
For each component, be sure to consider internal business operations, customers,
suppliers and partners. Keep in mind you are an E-Business with customers, suppliers,
and partners - you MAY NOT simply block everything!

You **must** include the complete policy (ACLs, ruleset, IPSec policy) in your paper. It is
not enough to simply state "I would include ingress and egress filtering…" etc. The
policies may be included in an Appendix if doing so will help the "flow" of the paper.

2.1 Security Policy

The security design principles of GIAC Enterprises have already been discussed under
Security Architecture. In that section we outlined the principle of defence in depth,
disallowing all traffic that is not explicitly allowed etc. In this section we will run
through how these policies are put into practice on the border router, firewall and VPN.

2.2 Border Router Configuration and Hardening

The router will be configured with two Ethernet cards in slots 0 and 1. Slot 0 will be the
external interface and Slot 1 the internal interface.

2.2.1 Ingress Filter

2.2.1.1   Access list is applied to the external Ethernet interface

```
! Ingress filter list 101 on Ethernet 0 interface
interface Ethernet 0
ip address 193.125.25.10        255.255.255.0
ip access-group 101 in
```

2.2.1.2   Deny Private Source Addresses

Private address ranges should never be used on the Internet, only on private networks
that use Network Address Translation (NAT) to communicate with the Internet.
Consequently any packets arriving at the router with private source addresses should be
dropped.

```
! block private address ranges (RFC1918)
```

Last printed 3/8/05 8:01 PM   Page 15 of 65

```
access-list 101 deny 10.0.0.0     0.255.255.255
access-list 101 deny 172.16.0.0  0.15.255.255
access-list 101 deny 192.168.0.0 0.0.255.255 any log
```

The 192.168.0.0 range is logged because these could be an attempt to attack the private
network ranges we are using internally.

2.2.1.3 Deny Localhost, Broadcast and Multicast addresses

Localhost (127.0.0.1) is used to test local IP configurations and should never be seen
"on the wire". Broadcast and multicast are not appropriate from the Internet because
they are not specific for our gateway

```
! block localhost, broadcast and multicast
access-list 101 deny 127.0.0.0        0.255.255.255
access-list 101 deny 255.0.0.0        0.255.255.255
access-list 101 deny 0.0.0.0          255.255.255.255
access-list 101 deny 255.255.255.255 0.0.0.0
access-list 101 deny 224.0.0.0        15.255.255.255
```

2.2.1.4 Deny reserved and unallocated addresses[v]

These address ranges have not been allocated by IANA or are DHCP or
autoconfiguration addresses.

```
! block reserved and unallocated addresses
access-list 101 deny 169.254.0.0 0.0.255.255    any log
access-list 101 deny 192.0.2.0   0.0.0.255      any log
access-list 101 deny 240.0.0.0   7.255.255.255  any log
access-list 101 deny 248.0.0.0   7.255.255.255  any log
```

2.2.1.5 Deny packets without a source address

A genuine communication would have a source address, otherwise the destination would
be unable to reply.

```
! Deny packets without a source IP address.
access-list 101 deny host 0.0.0.0 any log
```

2.2.1.6 Deny packets that use the address range of our network

Packets with these source addresses should be coming from within, not outside, our
network

```
! Deny inbound packets that use our source addresses.
access-list 101 deny NNN.NNN.NNN.0 0.0.0.255 any log
```

2.2.1.7 Permit remaining traffic to pass through the router to the firewall

```
! Permit remaining traffic
access-list 101 permit any
```

2.2.2 Egress Filter (list 102)

This access list is applied to the internal Ethernet interface

```
! Egress filter list 102 on Ethernet 1 interface
interface Ethernet 1
ip access-group 102 in
```

2.2.2.1 The only packets we will allow to leave our network are those with our network's source addresses

Any packets that aren't from our network's source address range will be dropped and logged, with the –input flag causing the MAC address of the host or previous router to be captured. This can tell us which host may have been compromised.

```
! Restrict egress to our network's address range
access-list 102 permit 193.125.25.0 0.0.0.255
access-list 102 deny any log –input
```

2.2.3 Restricting Administrative Access

We will limit administrative access to the router's virtual TTY ports to be via ssh from a specified host. The session will also be password protected and the login set to time out after 30 seconds to prevent denial of service attacks that simply sit at the login prompt of all the available vty sessions

```
! restrict admin access to 192.168.40.254
access-list 103 permit host 192.168.30.222
access-list 103 deny any

! apply to all five vtys
line vty 0 4

! restrict telnet, forcing secure shell instead
transport input ssh

access-class 103 in
login

password SECRET

! cause idle session to time out after 30s
exec-timeout 1 30
```

2.2.4 Runtime Environment Hardening[vi]

2.2.4.1 Login Banner

A login banner protects the company by making it clear that unauthorised access is not allowed and systems are monitored. This pre-empts the defence that no warning was given or personal privacy was compromised.

```
banner /
```

```
        WARNING: Unauthorised access prohibited.
        This system is actively monitored.
        /
```

2.2.4.2   Passwords

The **service password-encryption** command uses a simple, easily cracked, Vigenere algorithm to encrypt passwords and similar strings against casual browsing. The **enable secret** command protects the password that gives administrative access to the router, using the MD5 hashing algorithm.

```
        ! enable password encryption
        service password-encryption
        enable secret
```

2.2.4.3   Turn off management services

In a large network, hosts may be managed using Simple Network Management Protocol but this is not a requirement for our relatively simple GIAC Enterprises network. SNMP and the HTTP administration interface should be disabled by default, but it does no harm to make sure. We also wish to disable Cisco Discovery Protocol (CDP) and bootp.

```
        ! Turn off management services
        no ip http server
        no ip bootp server
        no snmp
        no cdp enable
```

2.2.4.4   Turn off other unused services

The TCP and UDP "small services" echo, chargen and discard are hardly ever used for legitimate reasons but can be used for denial of service attacks (e.g. using chargen to generate packets which are then reflected back by echo to cause a positive feedback loop).

```
        ! Turn off small services.
        no service tcp-small-servers
        no service udp-small-servers
```

Finger is used to determine the users logged into a device

```
        ! Turn off finger
        no service finger
```

Network Time Protocol (NTP) can be valuable to synchronise the clocks of various hosts in order to track attacks through different logs. However, with our own NTP server on the service network, we have no need to expose our service to the Internet and no need to set our clocks from an external service. Therefore NTP should not pass through the router.

```
        ! Turn off ntp
        no ntp enable
```

2.2.4.5 Turn off loose source routing

Loose source routing is used to redirect packets arriving at a remote system to another location and can be used to bypass access control lists.

```
! Disable loose source-routed packets.
no ip source-route
```

2.2.4.6 Limit ICMP

ICMP is used for IP network management, for example by telling a source host when a destination is unreachable. This may provide more information about our network than we wish to share with unknown third parties. It can also be used to give rise to denial of service attacks when the router fires off a storm of responses in answer to malicious directed broadcasts, such as smurf.

```
! Limit ICMP responses
no ip directed-broadcast
no ip unreachables
```

2.2.4.7 Protect router from overload

Floods can cause the router to be so busy responding to interrupts from the network interfaces that it is unable to do any other work. The **scheduler interval** and **scheduler allocate** commands ensure that the router stops responding to interrupts at regular intervals in order to service its work queue. Scheduler allocate is the more recent command, specifying the number of milliseconds to run with interrupts enabled (400ms - 60000 ms) and then with interrupts masked (100ms – 4000ms).

```
! Protect router from overload
! start in middle of range and fine tune if necessary
scheduler allocate 30000 2000
```

2.3 Firewall Configuration

Firewall configuration is considered in depth in the tutorial section. The rules are summarised here.

*   Allow SSH access from management console to border router
*   Allow management console access to hosts on third party network
*   Allow management console access to hosts on service network
*   Allow hosts on internal network access to mailserver on service network
*   Allow hosts on external network access to mailserver on service network
*   Allow hosts on external network access to web server on service network
*   Allow authenticated access from third party network to web server on service network
*   Allow authenticated access from third party network to internal network
*   Allow hosts on internal network access to ftp, http, ntp servers on service network
*   Allow proxy server on internal network, http and ftp access to the Internet
*   Deny any other access from internal hosts to the Internet

- Allow web server access to database server
- Deny everything else (default drop, log and alert rule)

2.4 VPN Architecture

Raptor's VPN service is configured by defining security policies. These policies are standard templates that are applied to individual connections – this makes configuration easier because settings do not have to be recreated for each tunnel.

Our Virtual Private Network service is required to provide two types of access; from mobile GIAC staff to the internal network and from customer companies to the secure web server. These two modes of operation have separate requirements so will be discussed separately and policies created for each.

Raptor offers three encapsulation types
   o   IPSEC/Static
   o   IPSEC/IKE
   o   swIPe

swIPe is proprietary and not supported by the PowerVPN server so is deprecated. Our choice is therefore between using static predefined algorithms or agreeing these dynamically when a connection is created using the Internet Key Exchange (IKE) protocol.

Raptor VPN is able to support the SHA1 and MD5 algorithms to authenticate the data payload and DES and 3DES to encrypt it. The security policy defines which of these algorithms will be used.

An IPSEC/Static connection will fail if either of the two parties is unable to support the specified protocols. This guarantees that either the session will be established with the required level of security or it won't be established at all.

IPSEC/IKE offers more flexibility by allowing the parties to negotiate for a common protocol, so for example a 3DES connection could fall back to DES if one of the parties wasn't able to support the stronger encryption. Raptor's IKE implementation also allows timeouts to be specified ensuring that tunnels aren't created from remote customers and then left open once they have finished their transactions.

Given that some of GIAC Enterprises' clients are likely to be in countries where encryption export restrictions apply we will use an IPSEC/IKE policy that permits both 3DES and DES for those connections that may be forced to use a weaker level of encryption, and a policy that specifies strong encryption for all cases where we know we are able to use it. This approach gives us flexibility because if a country that previously forced us to use weak encryption changes its rules to allow strong encryption (e.g. France) the remote end of the connection could have 3DES enabled without the necessity to change the policy on the local security gateway.

Conversely only our own staff will use mobile clients to access GIAC Enterprises' network, so for these we can define a static policy using strong encryption.

2.4.1 VPN Access By Mobile clients

Our mobile clients will use the RaptorMobile 6.5 software, which acts as both network entity and security gateway for the remote tunnel endpoint. These concepts are covered in detail in the tutorial on Raptor Firewall later in this document.

Clients will also use a personal firewall (Blackice) and antivirus software, which will be managed over the VPN connection using Microsoft System Management Server.

2.4.1.1   Security Policy For Mobile Clients

Parameter	Value	Comment
Name	IPSEC_static_mobile	
Description	IPSEC/Static policy for mobile clients	
Encapsulation Protocol	IPSEC/Static	
Pass traffic to Proxy services	True	NAT requires that traffic is passed up the stack
Data Integrity Algorithm	MD5	
Data Privacy Algorithm	3DES	
Tunnel Mode	True	Required for host – gateway communications
AH or ESP	ESP	Wish to encrypt, not just authenticate the packet. ESP is required for NAT.
Obsolete Protocol Version	False	Do not have any RFC1825 clients.

2.4.2 VPN Access From Customers' Networks

As we have said, the policy we apply to customer connections depends on whether they are able to use strong encryption. If they are then we can set up an IKE connection that is very similar to that for mobile clients, but with the addition of timeouts, otherwise we will use an IKE policy that sets the best available common standard.

Although IKE has the capability to fall back to <NONE> as the authentication or encryption algorithm we will always have a minimum standard that requires both authentication and encryption. We will never allow traffic to be passed in the clear.

2.4.2.1   Security Policy For Customers using strong encryption

Parameter	Value	Comment
Name	IPSEC_IKE_hisec	
Description	IPSEC/IKE policy for customers with 3DES	
Encapsulation Protocol	IPSEC/IKE	
Pass traffic to Proxy services	True	NAT requires that traffic is passed up the stack
Data Integrity Preferences	MD5, SHA1	Attempt MD5 first then SHA1

Data Privacy Preferences	3DES	Only allow 3DES
Data compression	<NONE>	Don't want to invest the CPU in compressing packets
Data Volume Timeout	100,000 kilobytes	Restrict the volume of data that can be downloaded in one session
Lifetime Timeout	480 minutes	Close the connection after 8 hours
Inactivity Timeout	15 minutes	Close connection if it is unused for 15 minutes
Tunnel Mode	True	Could use transport mode but only saves 20 bytes per packet at the cost of addressing flexibility
AH or ESP	ESP	Wish to encrypt, not just authenticate the packet. ESP is required for NAT.
Perfect Forward Secrecy	Enabled	Prevents key guessing
Diffie Hellman Preference	2 then 1	Use 1024 bit keys falling back to 768 bit

2.4.2.2   Security Policy For Customers forced to use weak encryption

Parameter	Value	Comment
Name	IPSEC_IKE_losec	
Description	IPSEC/IKE policy for customers with DES	
Encapsulation Protocol	IPSEC/IKE	
Pass traffic to Proxy services	True	NAT requires that traffic is passed up the stack
Data Integrity Preferences	MD5, SHA1	Attempt MD5 first then SHA1
Data Privacy Preferences	3DES, DES	Attempt 3DES first then DES
Data compression	<NONE>	Don't want to invest the CPU in compressing packets
Data Volume Timeout	100,000 kilobytes	Restrict the volume of data that can be downloaded in one session
Lifetime Timeout	480 minutes	Close the connection after 8 hours
Inactivity Timeout	15 minutes	Close connection if it is unused for 15 minutes
Tunnel Mode	True	Could use transport mode but only saves 20 bytes per packet at the cost of addressing flexibility
AH or ESP	ESP	Wish to encrypt, not just authenticate the packet. ESP is required for NAT.
Perfect Forward Secrecy	Enabled	Prevents key guessing
Diffie Hellman Preference	2 then 1	Use 1024 bit keys falling back to 768 bit

Part 2 – Security Policy Tutorial

Select **one** of the three security policies defined above and write a tutorial on how to implement the policy. Use screen shots, network traffic traces, firewall log information, and/or URLs to find further information as appropriate. Be certain to include the following:

1. A general explanation of the syntax or format of the ACL, filter, or rule for your device.

2. A general description of each of the parts of the ACL, filter, or rule.

3. An general explanation of how to apply a given ACL, filter, or rule.

4. For each ACL, filter, or rule in your security policy, describe:

   o the service or protocol addressed by the rule, and the reason this service might be considered a vulnerability.

   o Any relevant information about the behavior of the service or protocol on the network.

   o If the **order** of the rules is important, include an explanation of why certain rules must come before (or after) other rules.

5. Select three sample rules from your policy and explain how you would test each rule to make sure it has been applied and is working properly.

Be certain to point out any tips, tricks, or potential problems ("gotchas").

2.5 Security Policy Tutorial

In this section we are going to work through how to configure a Raptor firewall with the policies we outlined in the previous section. We will be using the Raptor Management Console (RMC) on our management client to remotely configure the Raptor firewall software on the firewall host. However some settings need to be made directly on the firewall.

Although this tutorial is not exhaustive it should provide enough detail to get the firewall system up and running and then create the rule set.

2.5.1 System Configuration

In order to be installed on Windows 2000, Raptor requires Windows 2000 Service Pack 2. The Windows 2000 installation should then be hardened, but OS hardening is beyond the scope of this tutorial. We will start from the point at which the firewall software is installed on the firewall server and the RMC is installed on the remote management client.

Some settings need to be made directly on the firewall server; the default gateway needs to be set, remote management needs to be enabled and static routes must be configured.

2.5.2 Setting Default Gateways And DNS
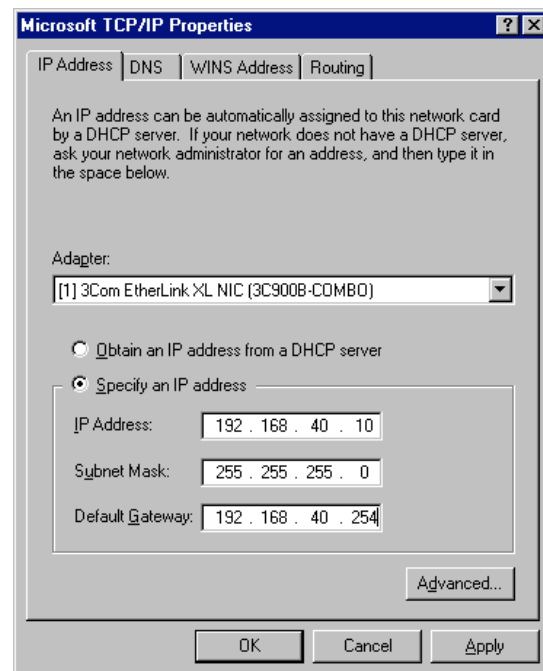
The default gateway of the **external** network interface card of the firewall should be set to point to the border router (IP 192.168.40.254)

The TCP/IP properties panel is accessed by clicking Start → Settings → Network and Dialup Settings

From the Network window, right click on the icon of the external interface and choose properties. The properties dialog for the NIC will then be displayed.

From the protocols list, select Internet Protocol (TCP/IP) and click the Properties button. This will cause the TCP/IP properties screen on the left to be displayed.

Enter the IP address of the interface, subnet mask and default gateway as shown. Then click the DNS tab and set the DNS service search order to start with the localhost address (IP 127.0.0.1) this will ensure that the DNS proxy on the firewall is used for name resolution.

N.B. Do not configure remote gateways for the internal interfaces of the firewall. These should be left empty but DNS pointed to localhost as above.

Hosts on the various GIAC Enterprises subnets must be configured with their default gateways and DNS pointing to the closest interface of the firewall; in this way all traffic beyond the local subnet will be directed to the firewall which will check and route it appropriately.

Host	Default Gateway and DNS service setting
Service Network hosts	192.168.20.10
Third Party Network hosts	192.168.10.10
Internal Network hosts	192.168.30.10

2.5.3 Enabling Remote Management

The Raptor Firewall must be configured to trust the Remote Management Console. This is achieved by running the **rempass** utility from the command line.

In the listing below, user responses are shown in **bold type.**

```
C:\Raptor\Firewall\Bin>rempass

 REMPASS - Host, password, service, and port configuration tool
-------------------------------------------
Enter one of the Rempass options shown below:

(A)dd new Host Configuration
(C)hange existing Host Configuration
(D)elete existing Host Configuration
(L)ist existing Rempass Host entries
(Q)uit Rempass

Rempass Option: A

Host name or IP address: 192.168.30.222
-------------------------------------------
Service List:

(1) Firewall Management Console
        -Configure firewall to accept remote management connections

(2) Logfile Retrieval
        -Configure firewall to allow remote client to access firewall logfiles

(3) Log Event Submission
        -Configure firewall to accept log output from remote client

(4) Content Scanning
        -Configure firewall to use remote content scanner

(5) Intrusion detection
        -Configure firewall to accept intrusion notification

Please Choose a Service ('m' for main menu): 1

Enter up to 64 characters for
192.168.30.222's passphrase: [passphrase]

Verify new password: [passphrase]
```

After the passphrase is accepted, quit from the rempass utility.

2.5.4 Configuring Static Routes

The firewall is only aware of the subnets to which it is directly connected. In order to route packets to the 192.168.50.0 secure network, which lies beyond the 192.168.30.0 internal network we must define a static route to the proximate interface of the internal firewall (IP 192.168.30.254) This is achieved from the command line:

```
route –p add 192.168.50.0 mask 255.255.255.0 192.168.30.254
```
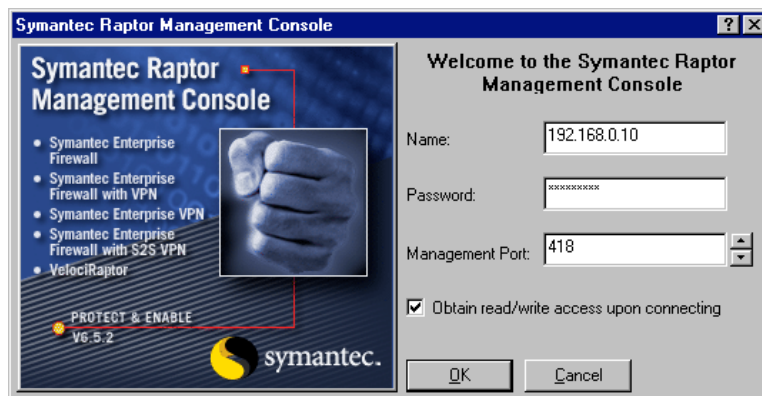
2.5.5 Raptor Management Console

From this point onwards all firewall management can be performed from the remote management console. It is started by double clicking its desktop icon or from Start → Programs → Symantec Raptor Management Console → Raptor Management Console.

The RMC is a plug in to Microsoft's Management Console so has the standard MMC user interface of a scope pane on the left with a result pane on the right.
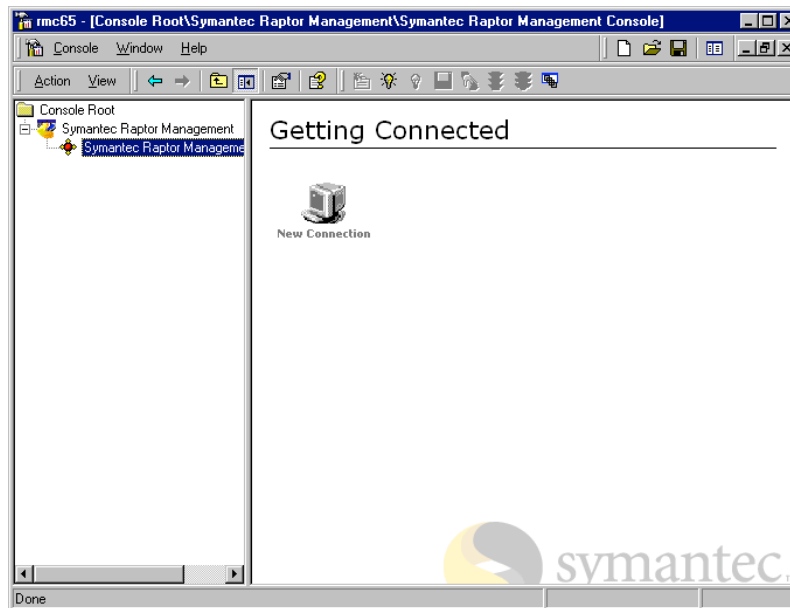
Items are selected from the directory tree in the scope pane and context sensitive property pages are displayed in the results pane.

The first time the RMC is run it displays the Getting Connected results page with one icon; New Connection. Double click this and the Console Creation Window appears.
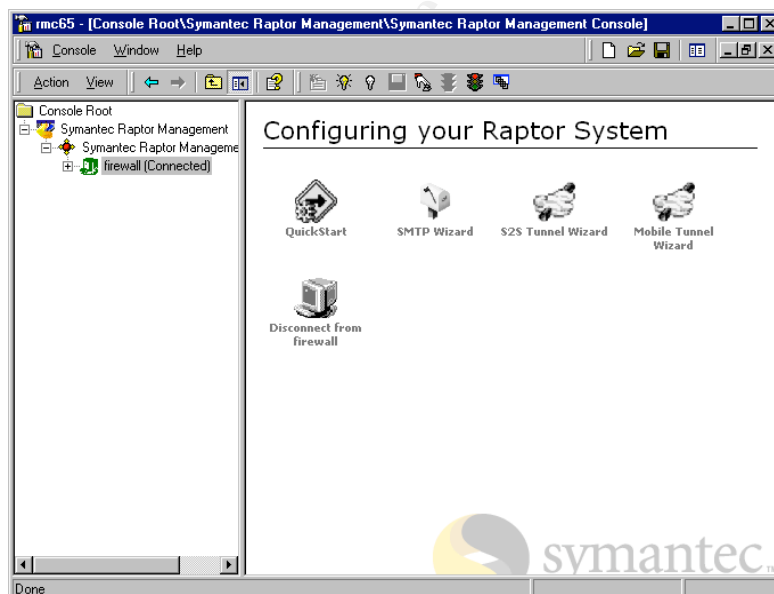
Enter the IP address of the firewall – that is the IP address of the firewall interface closest to the management console.

Then enter the passphrase and click OK.

This will establish a connection with the firewall.
The RMC then displays the Root Directory Window which has a number of icons for quick configuration wizards.

We can confirm that we have a read/write connection to the firewall because its icon in the scope pane is green. A read only connection would be grey.

This screenshot shows the scope pane with the management directory tree fully expanded. The next stage of configuring the system is to work through the Base Components, configuring each in turn.
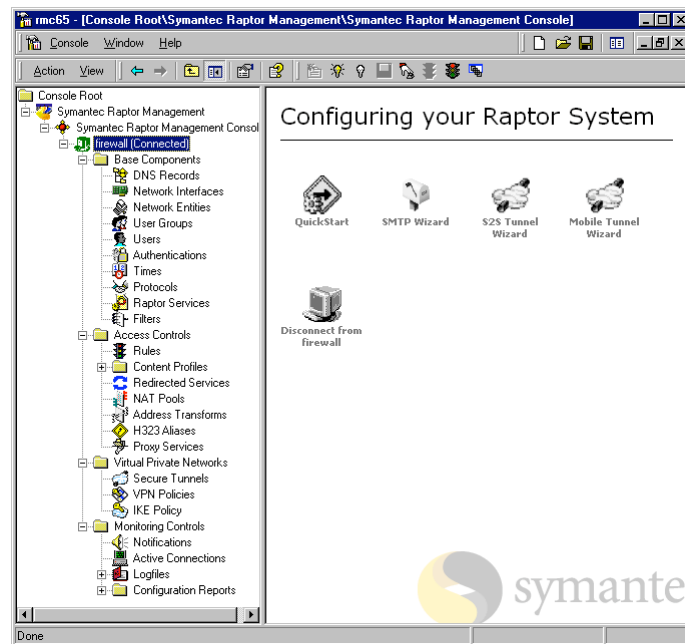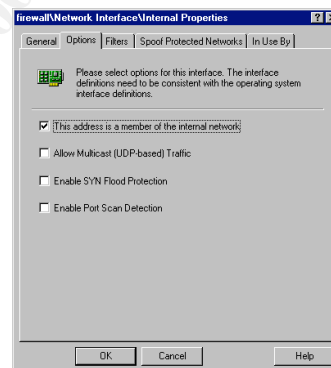

2.5.6 Defining Network Interfaces

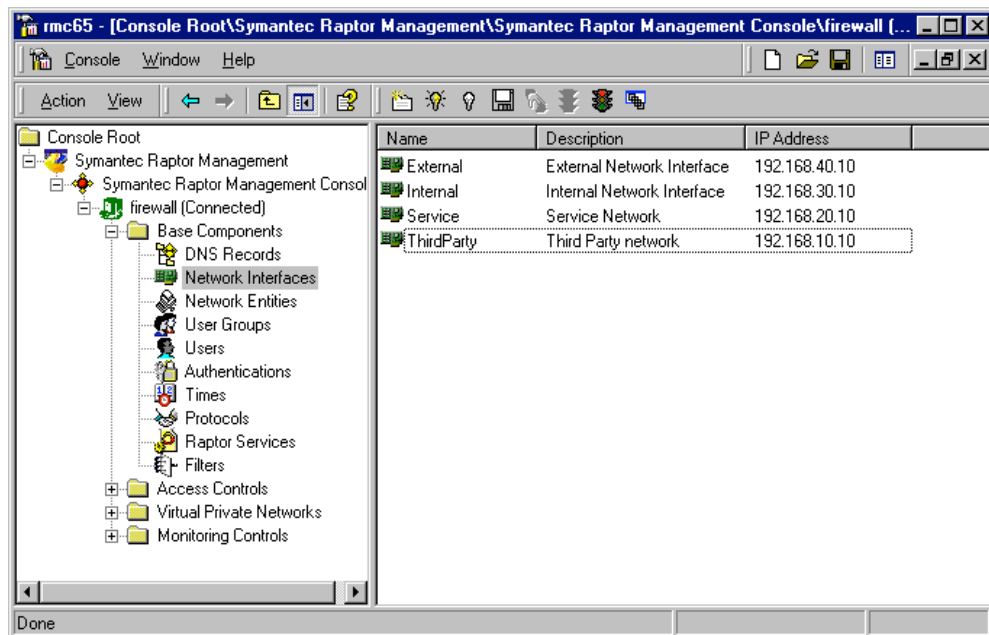Right click on the network interfaces leaf and select New → Network Interface
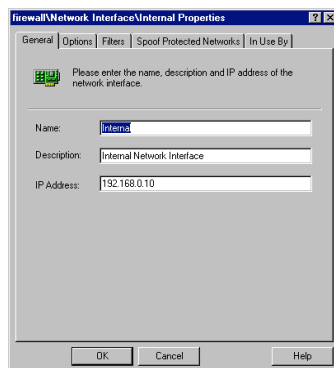
For each network interface enter its name, description and IP address as shown in the dialog on the left.

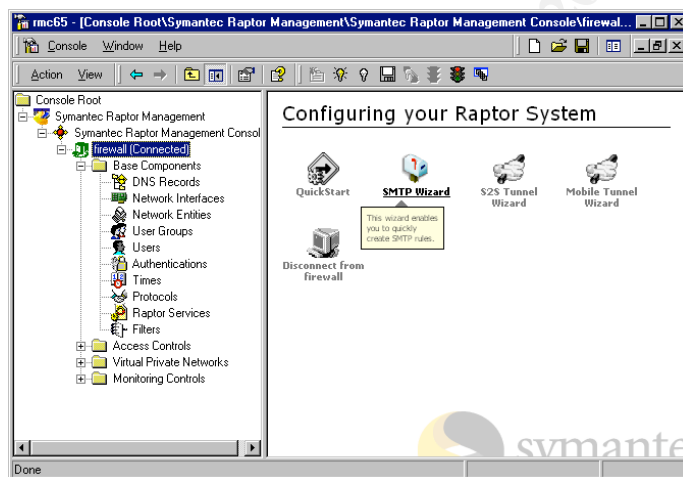If the interface is Internal, check the "address is a member of the internal network" checkbox.

Once all interfaces have been defined they will appear as shown in the screenshot below.
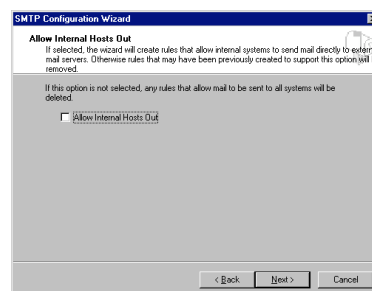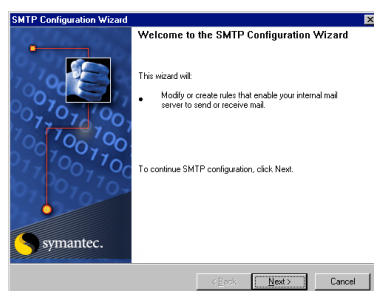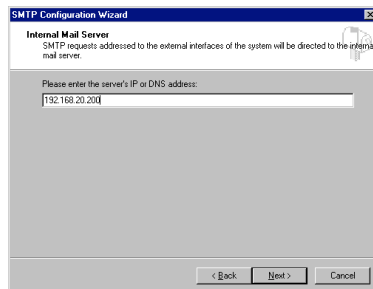

2.5.7 SMTP Configuration Wizard

Raptor has a number of configuration wizards to set up more complex services, such as simple mail transport protocol.

To start the wizard click on the SMTP Wizard icon in the Root Directory window.

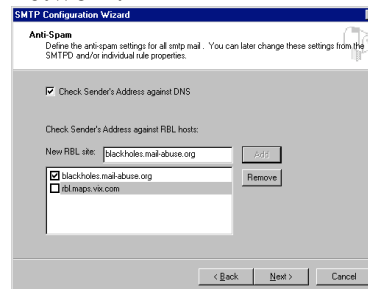Click Next >   We do not want our users accessing SMTP
                                                servers on the Internet so ensure the check
                                                box is clear. This forces all SMTP traffic
                                                to go via the SMTP server on the service
                                                network.

Enter the IP address of the SMTP server.   Anti-spam measures can be implemented
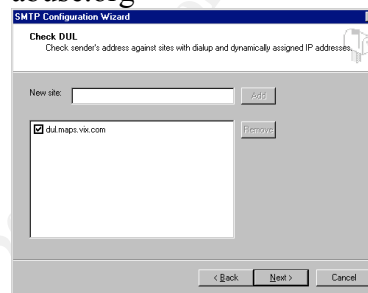   by checking against a list of known
                                                spamming hosts.

                                                N.B. rbl.maps.vix.com is no longer live
                                                and directs users to blackholes.mail-
                                                abuse.org
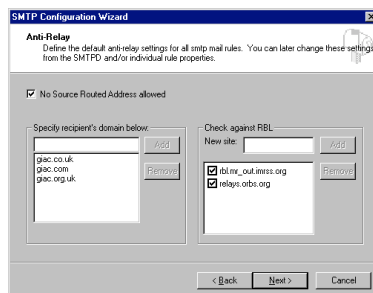
This dialog defines the domains for   Check whether mail has originated from
which we will accept incoming mail and   Dial Up users trespassing on another
the databases we will check to see if   domain's SMTP server.[vii]
mail has originated from a known
spamming relay.

Rules for the firewall are then created which should be manually checked and
modified if necessary.

2.5.8 Configuring the DNS Daemon

Raptor can either reference separate DNS servers or maintain its own DNS on the
firewall server. For ease of administration we have chosen to run the DNS on the
firewall server but our network design allows for an external DNS server on the service
network and an internal one on the internal network should name resolution traffic
impact the performance of the firewall.

The Raptor DNS service allows hosts to be specified as internal (Private) or external
(Public) with only those designated Public being resolved in response to requests from
outside the GIAC Enterprises domain. This provides the same functionality as a split
DNS.

Last printed 3/8/05 8:01 PM   Page 29 of 65

Right click on DNS records in the scope pane, then select New Host. The dialog shown on the left appears.

Set whether the host is Private or Public, its type, fully qualified domain name, IP address, alias(es), description and the domains it serves.

Repeat for each host on the network.

When the DNS is fully populated it will appear as shown below
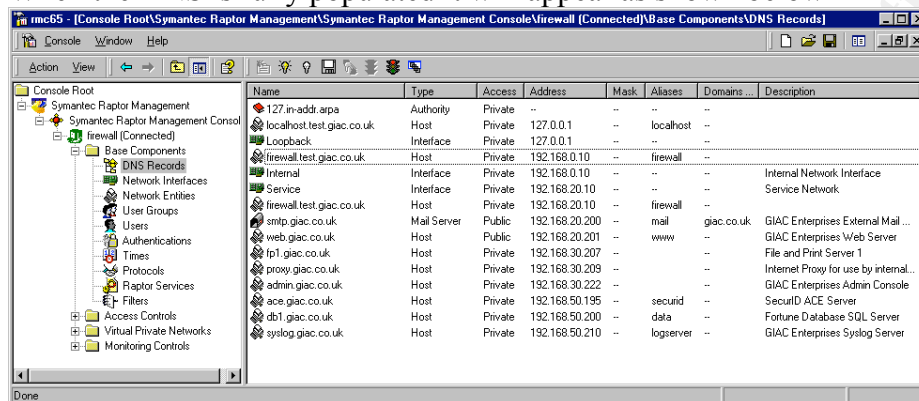
2.5.9 Defining Network Entities

Network entities are hosts, groups of hosts, subnets and domains. Defining these items as entities provides a convenient label to refer to them when creating rules.
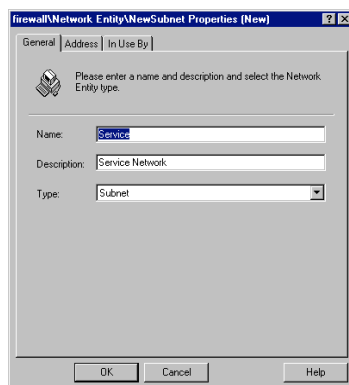
Right click on Network Entities in the scope pane and select New and the type of entity to be created.

In this case we are creating a new subnet entity so specify its name, description and type.
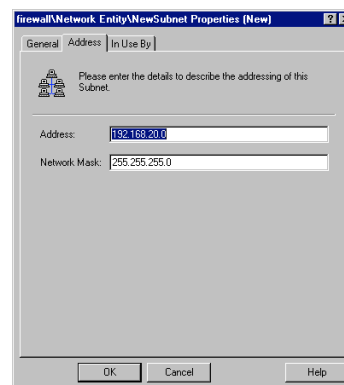
The address tab then lets us define the subnet address and mask.

The In Use By tab isn't used at this stage but is used later to view and modify rules involving this entity.

When all network entities have been defined they will appear as below
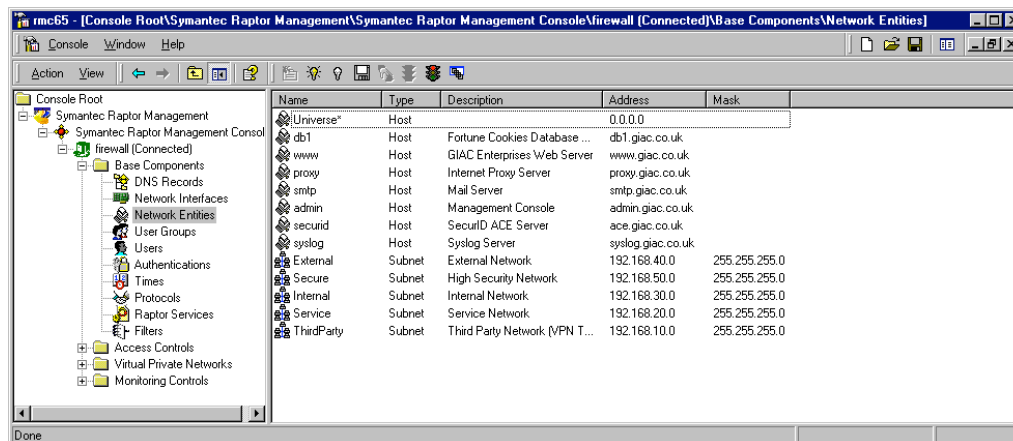
2.5.10   Referencing The SecurID ACE Server

Configuring the ACE Server is beyond the scope of this tutorial, but the firewall needs to be aware that the server exists and where to point requests to it. This is controlled from the Authentications dialog.

Click the Authentications leaf in the scope pane to display the possible authentication mechanisms. This is shown in the dialog below.
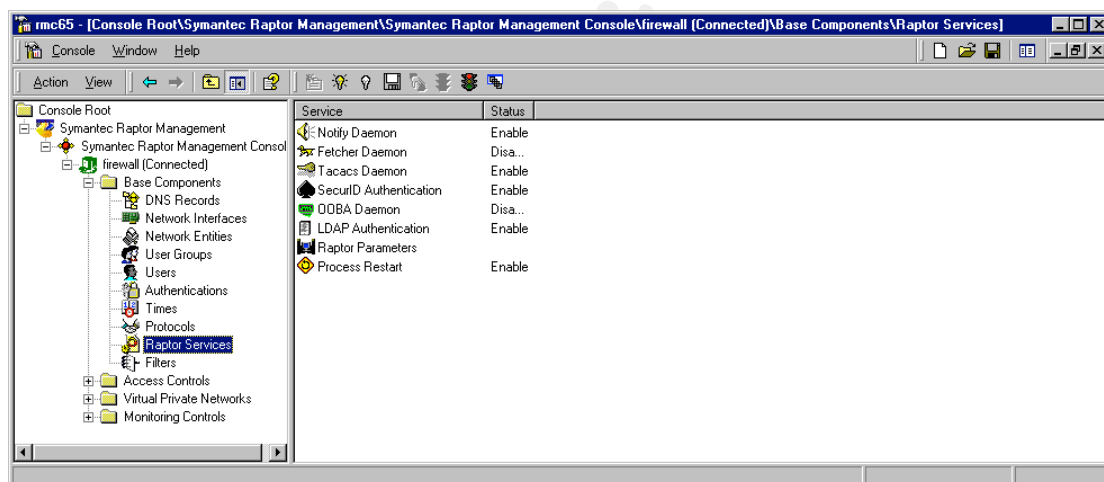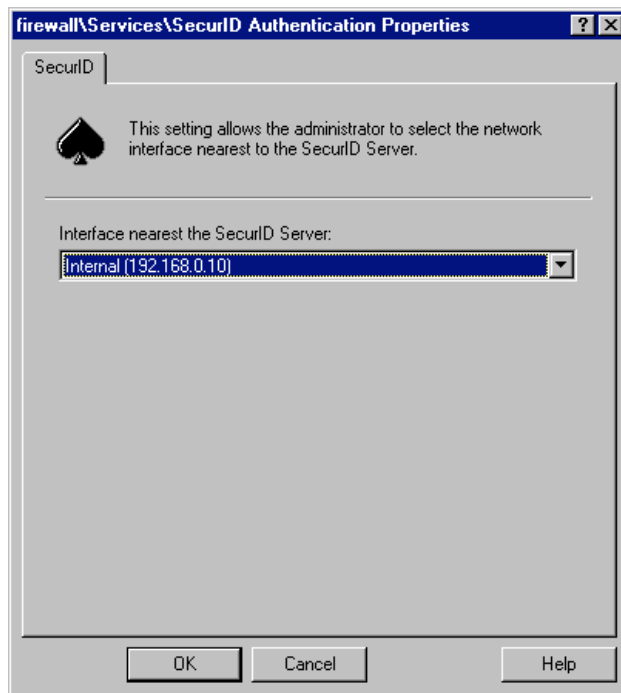
Double click the SecurID
Authentication service to
display the configuration
dialog shown on the left.
Then specify the interface
closest to ACE server. As
the ACE server is on the
secure network, the closest
interface will be the internal
network.

2.5.11   Creating Access Controls

Access rules are configured by expanding the Access Controls folder and clicking on
rules.
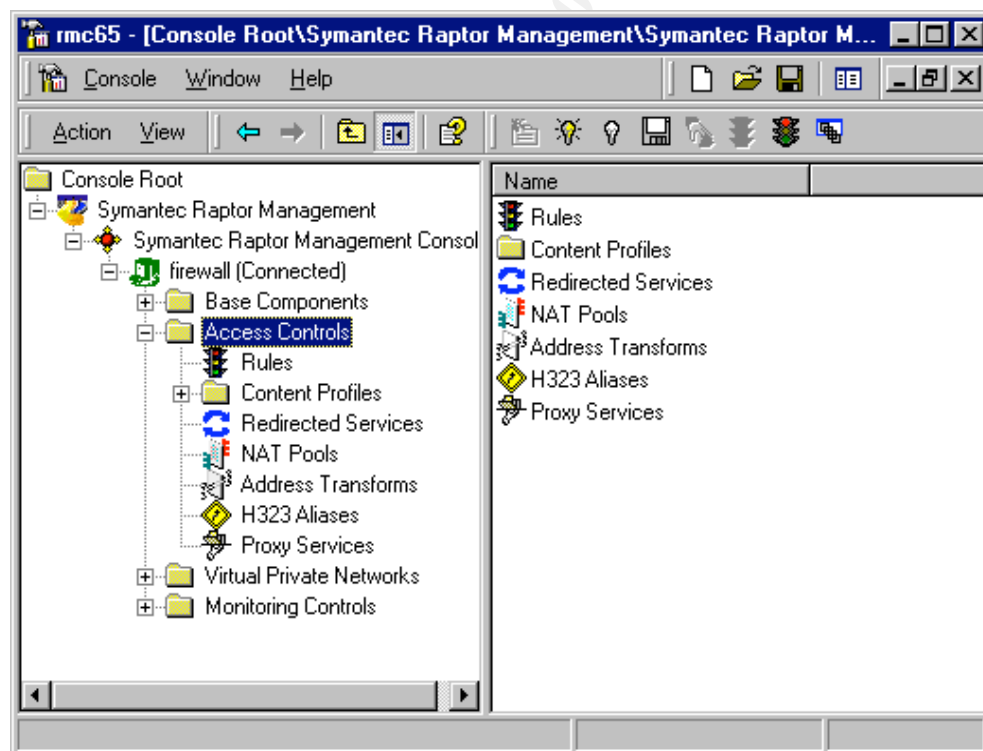
Each rule is then configured using the dialog below
.
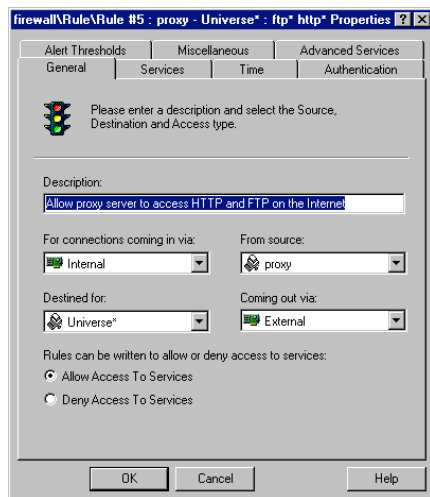
The General tab allows a description to be given of the rule, the interface through which traffic comes in and out of the firewall server and the source and destination network entities.
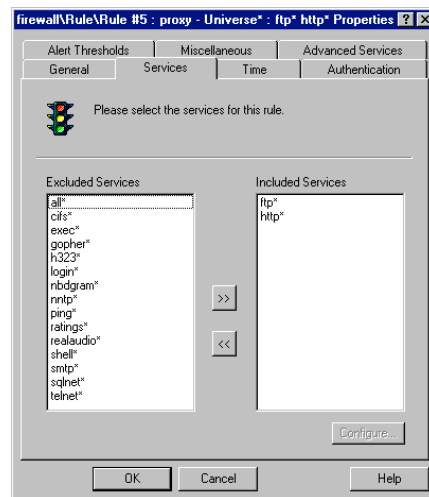
The services tab enables you to specify the services to be allowed (or blocked in the case of a blocking rule).

The time tab allows the time to be set for which the rule is active. Most of our rules will be set to be permanently active: <ANYTIME>
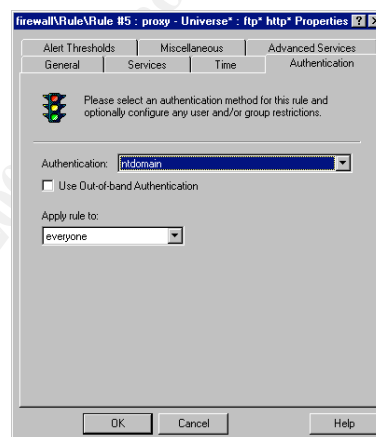
The authentication tab allows you to specify the authentication scheme used with the rule.
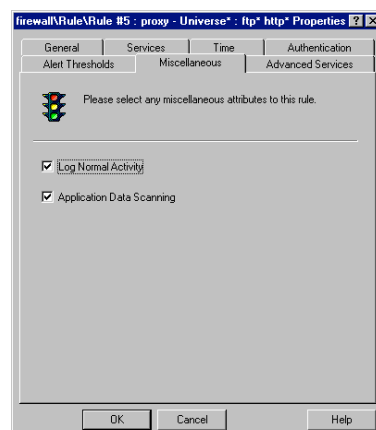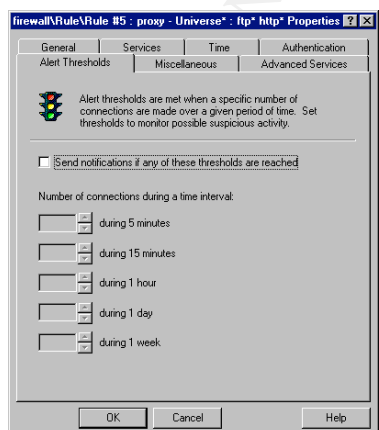
Alerts may be set if the rule is invoked more times in a given period than the specified threshold. This is particularly useful if rules are set for intrusion detection – e.g. traffic to sensitive hosts out of business hours.

The miscellaneous tab allows logging and application data scanning to be set.

Advanced services allows protocols and services to be configured which aren't part of the built-in list

2.6 Rule Set For GIAC Enterprises' Security Policy

The following rules will be configured on the GIAC firewall, in the order specified.

2.6.1 Administrative Rules

Administering the firewall itself: No rule is set to control administrative access to the firewall. Instead this is controlled using the **rempass** utility to create a trust relationship between the Management Console and firewall server, as previously described.

Administering perimeter hosts: Our security policy states that all hosts on the external, third party and service networks that are remotely managed will be managed from one management console. We will set a separate rule to allow access from the management console to each of these subdomains. This provides better granularity than a single rule that allowed the management console access al all interfaces.

SecurID is used on to provide two factor authentication with these rules and a low alert threshold will be set to ensure that any activity is noticed and reconciled with the change log. Time restrictions will not be set because administrative changes may need to be made at any time of the day or night.

There are no rules to manage domain transfers and other DNS activity because this is handled by the firewall itself, rather than such traffic routing through the firewall. There is also no rule to drop noisy traffic because we don't expect our logs to be too large at first and want to capture all data for analysis.

Rule Name:	Rule #1: admin – ExtRouter: shell		
Description:	Allow Management Console to administer router (SSH only)		
In Via:	Internal	**Source:**	Admin

Out Via:	External	Destination:	ExtRouter
Permissions:	ALLOW	Services:	Shell
Time:	<ANYTIME>	Authentication:	Securid
Logging:	Yes	Application Data:	No
Comments:	The border router is our most exposed perimeter device. With this rule we specify that it may only be managed from one specific management console, using SSH to protect passwords and session data and that SecurID two factor authentication must be used for the traffic to traverse the firewall. Optionally we could specify the MAC address of the NIC on the management console to protect against address spoofing, but that could backfire if the network card needed to be replaced. As the router is the only device on this subnet we are setting the rule with the router as the destination rather than the subnet.		

Rule Name:	Rule #2: admin – ThirdParty: all		
Description:	Allow Management Console to administer hosts on third party network		
In Via:	Internal	Source:	Admin
Out Via:	ThirdParty	Destination:	ThirdParty
Permissions:	ALLOW	Services:	All
Time:	<ANYTIME>	Authentication:	Securid
Logging:	Yes	Application Data:	No
Comments:	Our security policy says that all third party connections will terminate on this subnet. This means we may have a variable number of termination devices and protocols. This rule allows the management console to access and configure these hosts.		

Rule Name:	Rule #3: admin – Service: all		
Description:	Allow Management Console to administer service network		
In Via:	Internal	Source:	Admin
Out Via:	Service	Destination:	Service
Permissions:	ALLOW	Services:	All
Time:	<ANYTIME>	Authentication:	Securid
Logging:	Yes	Application Data:	No
Comments:	Similar to the previous rule, this allows us to administer hosts on the service network – again with two factor authentication and logging.		

2.6.2 SMTP Rules

We want to ensure that our internal hosts only send mail via the mail server on the server on the Service network. This server also holds their POP3 mailboxes, which is a vulnerability because the service network is less trusted than the Internal or Secure networks. An upgrade path would be to set up an internal mail server that regularly polls

the SMTP/POP3 server on the external network and point clients at that instead. We also recommend anti-virus and content scanning capabilities be installed on the SMTP server or a separate server on the service network to allow email to be checked as it enters and leaves GIAC Enterprises.

Rule Name:	Rule #4: universe – smtp: smtp pop3 imap4		
Description:	Allow hosts on inside network to send mail to mailserver		
In Via:	Internal	Source:	Universe
Out Via:	Service	Destination:	smtp
Permissions:	ALLOW	Services:	Smtp, pop3 imap4
Time:	<ANYTIME>	Authentication:	<NONE>
Logging:	Yes	Application Data:	No
Comments:	This is a fairly relaxed rule that allows any hosts on the internal network to send internet mail without authentication. Traffic is logged through. Consideration could be given to using NT authentication to provide user as well as host logging but this could preclude applications sending messages (e.g. IDS to pager or SMS message gateways) it they do not run in a user context.		

Rule Name:	Rule #5: universe – smtp: smtp		
Description:	Allow external clients to send mail to mailserver		
In Via:	External	Source:	Universe
Out Via:	Service	Destination:	smtp
Permissions:	ALLOW	Services:	smtp
Time:	<ANYTIME>	Authentication:	<NONE>
Logging:	Yes	Application Data:	No
Comments:	This rule allows us to receive mail from external parties. We maintain separate rules for internal clients and external ones for ease of management and also because making the source <ANY> would allow smtp connections from the third party network and we don't want our partners directly accessing our mail server. Also we don't want external parties using POP3 or IMAP4 Although not precluded by this rule, we are aware of the danger of being used as a mail relay and would configure our mail server not to allow it.		

2.6.3 HTTP Rules

We have three directions in which to route traffic to the GIAC web server:
- The public and Suppliers must be able to access the GIAC web site from the Internet
- Customers and Partners must be able to access the GIAC web site from the third party network
- Staff must be able to access the GIAC web site from the internal network

In addition staff need to be able to access the Internet, via an HTTP/FTP proxy server.

Rule Name:	Rule #6: universe – www: http		
Description:	Allow external clients to access GIAC web server		
In Via:	External	Source:	Universe
Out Via:	Service	Destination:	www
Permissions:	ALLOW	Services:	http
Time:	<ANYTIME>	Authentication:	<NONE>
Logging:	Yes	Application Data:	Yes
Comments:	This rule allows any external hosts to access the GIAC web server without authentication. The general public will be able to browse general company information and suppliers authenticate with the server to enter fortune cookie sayings. Activity is logged in the firewall and web server logs and application data scanning enabled to protect against trojans and script attacks.		

Rule Name:	Rule #7: Universe – www: http		
Description:	Allow connections from third party network to GIAC web		
In Via:	ThirdParty	Source:	Universe
Out Via:	Service	Destination:	www
Permissions:	ALLOW	Services:	http
Time:	<ANYTIME>	Authentication:	Securid
Logging:	Yes	Application Data:	Yes
Comments:	This rule allows an authenticated connection from the third party network to the web, giving access to secure pages for high-value transactions such as downloading collections of fortune cookie sayings.		

Rule Name:	Rule #8: Universe – any		
Description:	Allow connections from third party network to Internal network		
In Via:	ThirdParty	Source:	Universe
Out Via:	Internal	Destination:	Internal
Permissions:	ALLOW	Services:	any
Time:	<ANYTIME>	Authentication:	Securid
Logging:	Yes	Application Data:	Yes
Comments:	This rule allows an authenticated connection from the third party network to the internal network. The SecurID authentication will only allow this access for GIAC staff who are members of a remote access group.		

Rule Name:	Rule #9: internal – service: ftp http ntp		
Description:	Allow internal systems to access hosts on service network		
In Via:	Internal	Source:	Internal
Out Via:	Service	Destination:	Service
Permissions:	ALLOW	Services:	ftp http ntp
Time:	<ANYTIME>	Authentication:	Ntdomain

Logging:	Yes	Application Data:	Yes
Comments:	This is quite a significant rule because it is doing several things: • It allows all internal hosts to access all servers on the service network • It explicitly allows all protocols used by the services on the service network but we don't set <ANY> protocols because this could allow a hostile insider to exploit a protocol we have failed to lock down at the server. • All connections require the user to have an NT domain (which is why we don't include smtp in this rule because that is unauthenticated in rule #4) • It allows non-proxied http connections from internal browsers – which would be set to not proxy the giac.co.uk domain. • NB hosts on the secure network will inherit permissions of the internal network once they pass through the internal firewall		

Rule Name:	Rule #10: proxy – Universe: http ftp		
Description:	Allow proxy server to access HTTP and FTP on the Internet		
In Via:	Internal	Source:	proxy
Out Via:	External	Destination:	Universe
Permissions:	ALLOW	Services:	http ftp
Time:	<ANYTIME>	Authentication:	<NONE>
Logging:	No	Application Data:	No
Comments:	This rule supports web browsing by users on our internal network providing it is done via the proxy server. We aren't logging at the firewall because the proxy logs will provide much better detail of our users browsing habits and we aren't authenticating because they will have authenticated to the proxy using NT challenge-response. At present we are allowing the users to browse at any time but if the line utilisation becomes too great we could restrict this rule to off-peak hours.		

Rule Name:	Rule #11: Internal – Universe: any		
Description:	Deny direct connections from internal clients to the Internet		
In Via:	Internal	Source:	Internal
Out Via:	External	Destination:	External
Permissions:	**DENY**	Services:	Any
Time:	<ANYTIME>	Authentication:	<NONE>
Logging:	Yes	Application Data:	No

Comments:	This rule is to prevent our users bypassing the Internet proxy and connecting directly. Such connections will be dropped and logged, so that we can educate them to follow security policy. This rule also protects against trojans that rely on a default route to the Internet, although many now use the browser settings instead of looking for a default gateway. Alerts can be set to detect traffic originating from the internal network.

Rule Name:	Rule #12: www – db1: sqlnet		
Description:	Allow web server to refer back to database server		
In Via:	Service	Source:	www
Out Via:	Internal	Destination:	db1
Permissions:	ALLOW	Services:	Sqlnet
Time:	<ANYTIME>	Authentication:	None
Logging:	Yes	Application Data:	No
Comments:	This rule breaks our default policy of only allowing clients to initiate communications with servers on less trusted networks. In this case the web server on the service network is connecting back to the database server on the secure network. To mitigate this risk we only enable sqlnet and apply additional monitoring on the internal firewall.		

Rule Name:	Rule #13 Universe – Universe: all		
Description:	Default drop, log and alert rule		
In Via:	<ANY>	Source:	Universe
Out Via:	<ANY>	Destination:	Universe
Permissions:	**DENY**	Services:	all
Time:	<ANYTIME>	Authentication:	Securid
Logging:	Yes	Application Data:	No
Comments:	This is the catch-all rule to manage any traffic that hasn't been picked up by a specific rule earlier. Its purpose is to allow better logging and alerting than if the firewall were allowed to silently drop packets that did not fit an allow rule.		

2.7 Test Three Rules

In this section we are required to test three of the rules defined above, both to demonstrate that the intended protocols and routing paths are allowed and also that unintended ones are blocked. Because the firewall is limiting connections to specific protocols we need a tool that will generate the appropriate packets – for example, it will not suffice to try to "ping" the remote system if ICMP is blocked.

To perform our tests we will use a laptop configured to dual-boot under Windows 2000 and Linux. This will give us a versatile client that can be moved between our different

subnets to test connectivity. Tools of choice include nmap[viii] under Linux and Retina[ix] and ws_ping[x] under Windows 2000.

We will use a second Windows 2000 laptop, running Netmon in promiscuous mode, to capture and analyse network traffic. We will also analyse firewall logs to confirm that disallowed packets are being dropped.

2.7.1 Allow Management Console to administer router (SSH only)

We previously said that the external router is our most exposed perimeter device and for this reason wish to limit administrative access to one host (the management console, 192.168.30.222) and one protocol, SSH.

Starting on the external network we can attempt to connect an SSH client to the router. However as the laptop would need an IP address on the external subnet (192.168.40.0) the connection should be rejected. Attempts to spoof the source IP should also fail because the laptop would not receive any replies.

Next we should repeat our test with the test laptop on the third party and service network. The firewall should not pass any traffic from the third party network to the router at all. Similarly no traffic should be initiated from the service network.

Moving the laptop to the internal network we can confirm that only the IP address of the proxy server, running http and ftp, and the management console running ssh should be able to send traffic through the firewall to its external interface. Spoofing addresses and protocols in the 192.168.30.0/24 range should confirm that the only packets reaching the external subnet are ssh from 192.168.30.222 and http/ftp from 192.168.30.209.

2.7.2 Allow external clients to access GIAC web server

Our next priority is to ensure that external parties can access the GIAC web server because this is the public face of the company. This rule is very easily tested, just put our laptop on the external subnet, run up a browser and enter the IP address. We should get straight to our web site. Once we have achieved this we can try accessing the web site across the Internet to confirm that we reach the site through the border router.

We should also confirm that http and sntp are the only protocols allowed though to the service network from the external network segment. We can confirm this by scanning the web and smtp server from the external subnet using nmap.

2.7.3 Allow connections from third party network to Internal network

Both our mobile staff and our customers utilise VPN connections to our third party network. Mobile staff can then access any host on the internal network using any protocol while customers can only access the web server on the service network using http. Both types of connection are authenticated using SecurID so it is imperative that all sessions initiated from the third party network are authenticated and that staff and customer groups are routed correctly.

Again we will start by placing our test laptop on the third party network and trying to map the third party and internal networks, without supplying any credentials. Next we

will authenticate as a member of the SecurID customers group and try to map the third
party and internal networks. This will demonstrate that only http is passed.

Finally we will authenticate as a member of the GIAC staff group and should then be
able to reach any host on the internal network with any protocol. We should not be able
to reach hosts on the service network.

3 Assignment 3 – Audit Your Security Architecture

You have been asked to conduct a technical audit of the **primary firewall** (described in Assignments 1 and 2) for GIAC Enterprises. In order to conduct the audit, you will need to:

1. Plan the audit. Describe the technical approach you recommend to assess the firewall. Be certain to include considerations such as what shift or day you would do the assessment. Estimate costs and level of effort. Identify risks and considerations.

2. Conduct the audit. Using the approach you described, validate that the primary firewall is actually implementing GIAC Enterprises' security policy. Be certain to state exactly how you do this, including the tools and commands used. Include screen shots in your report if possible.

3. Evaluate the audit. Based on your assessment (and referring to data from your assessment), analyze the perimeter defense and make recommendations for improvements or alternate architectures. Diagrams are strongly recommended for this part of the assignment.

Note: DO NOT simply submit the output of nmap or a similar tool here. It is fine to use any assessment tool you choose, but you must annotate/explain the output.

3.1 Plan The Audit

3.1.1 Authorisations and Timing

It is essential that senior management authorisation has been obtained before conducting any kind of security audit or penetration test. Senior management should be made aware of the scope of the test and the possibility that it may impact business systems. They should advise on the best times to perform tests in order to mitigate this risk.

In the case of GIAC Enterprises, we are informed that Internet activity is greatest in the early hours, because most business partners are located in the Far East. We can confirm from network bandwidth utilisation that incoming traffic reaches a peak around 10am GMT and tails off by 1pm. GIAC staff tend to use the afternoon to authorise new sayings and for financial updates, leaving work by 6pm.

We therefore plan our internal auditing activities to run from 6pm to midnight, leaving a two hour contingency from midnight to 2am when traffic typically starts to increase.

3.1.2 Personnel

We agree with senior management that an independent auditing company should be retained to perform and document the audit. This ensures that a fresh look will be taken at our implementation, which may uncover vulnerabilities that we have missed. All work will be governed by a confidentiality and non-disclosure agreement.

3.1.3 Scope

Statements of work are prepared by the auditors, and agreed with senior management, for each of the following activities.

3.1.3.1   External Perimeter Security Audit

In this procedure the auditors work from outside our perimeter to see what information they can uncover. This could include DNS and Whois information, port scans, war dialing against published phone numbers and social engineering. These are all techniques a hacker might use to gain unauthorised access[xi].

At this stage we specifically exclude denial of service attacks because we wish to address availability and business continuity separately once the recommendations of the security auditors have been implemented. That is, we want to make sure our system is fully hardened before trying to knock it over. We also ask our auditors to restrict their activities to the primary firewall rather than any other external facing systems they may identify.

The deliverable of the External Perimeter Security Audit will be a report detailing what information the auditors were able to uncover, any theoretical vulnerabilities, the results of any vulnerabilities they were able to exploit in order to compromise our firewall and recommendations, and cost, to fix the vulnerabilities.

3.1.3.2   Architectural Review

Once the auditors have had the opportunity to scan our perimeter for themselves and report on what they could find, we will give them the complete answer in the form of our Security Architecture document.

The deliverable of the Architectural Review will be a report reviewing our security architecture against industry best practice. Any theoretical vulnerabilities will be highlighted together with recommendations, and cost, to fix them.

At this stage the auditors may wish to return to the External Perimeter Security Audit to demonstrate the vulnerabilities. We will allow this because the fact that they didn't uncover the vulnerability in the first stage doesn't mean that a hacker wouldn't – we have no wish to rely on security through obscurity.

3.1.3.3   Internal Network Security Audit

Next we will give the auditors access to our Internal network so that they can attack the firewall from the same position as a trusted employee. They will now be within our security perimeter so we will be asking them to report on scope for sabotage and fraud by disgruntled employees.
The deliverable of this stage will be a report listing any vulnerabilities of the internal firewall interface both from an infrastructure point of view and a human resources point of view.

3.1.3.4   Audit of Firewall Configuration

Next we will ask our auditors to review the configurations of the firewall. In this step we wish to confirm that it has been configured as planned, that all necessary patches have been applied, that unused ports are closed etc.

The deliverable from this stage will be a report of the compliance of the firewall with the documented security policy, along with recommendations of how to bring non-compliant features into line and recommendations of additional hardening steps that had been missed in the initial policy.

3.1.3.5 Process Audit

Finally we will ask the auditors to review the workflow and processes by which staff and partners are granted access to systems, and firewall configuration is maintained. We will ask them to review the application and authorisation procedures for external partners to be granted access through the firewall, induction and security awareness training for staff accessing the Internet through it, change control and any other process steps by which a vulnerability may be introduced to our security infrastructure.

The deliverable of this stage will be a report documenting our workflow and areas in which it could be subverted for a malicious party to get access to our systems, or a mistake made that could introduce a vulnerability that could later be exploited, or result in unavailability of the service.

3.1.4 Risk Analysis

Having conducted the audits we will be presented with a list of hard (infrastructure) and soft (people/processes) vulnerabilities – probably far more than we can hope to fix in the short term. A risk analysis would allow us to decide where to place resources in order to achieve the greatest mitigation for our investment.

A vulnerability, alone, does not equate to risk. We do not know how likely it is someone will try to exploit the vulnerability (the Threat) or the damage the vulnerability will do if it is exploited (the Impact). Ideally we would seek to apply values to each of these parameters to give an overall risk, in terms of cost:

Threat x Vulnerability x Impact = Risk

Applying this formula to each or our vulnerabilities gives a measure of the financial cost of the risk.

The audit should have identified mitigation steps for the vulnerabilities with an indication of cost, and it is likely that some mitigation steps will fix several vulnerabilities, so the next step would be to look at the steps at our disposal and decide which of these will reduce the greatest value of risk.

Our final audit report to management should include this summary allowing them to decide how much money they are prepared to invest in mitigating IT security risk as it relates to the firewall.

3.2 Conduct The Audit

3.2.1 External Perimeter Security Review

3.2.1.1   DNS Information
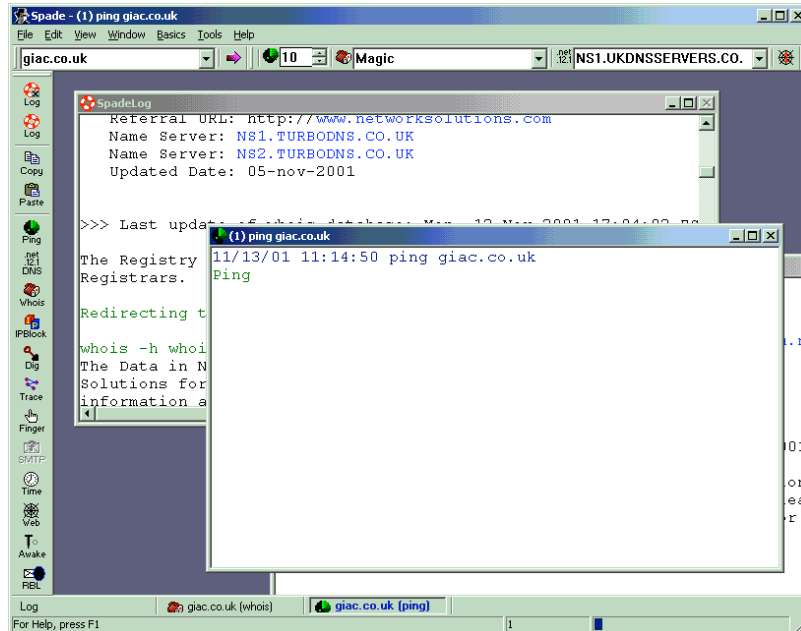
Our first stage is to try to work out what traffic is routing through the firewall. A good start is to see what we can find out about the giac.co.uk domain. The tools at SamSpade.org provide an excellent starting point because they allow us to run our searches both from our local network, using the Windows client, and from the Samspade.org web site.

This is what we were able to find out about **giac.co.uk**

whois giac.co.uk

```
11/10/01 10:49:44 whois giac.co.uk
.uk is a domain of United Kingdom
(international dialing code 44)
Searches for .uk can be run at http://www.britain.eu.net/naming-co/whois-form.html

whois -h whois.nic.uk giac.co.uk ...

   Domain Name: GIAC.CO.UK
   Registered For: GIAC Enterprises Ltd
   Domain servers listed in order:
   GATE.GIAC.CO.UK                  193.125.25.10
   NS.PIPEX.NET                     158.43.128.26

   WHOIS database last updated at 02:35:01 10-Nov-2001

The NIC.UK Registration Host contains ONLY information for domains
within co.uk, org.uk, net.uk, ltd.uk and plc.uk.  Please use the whois
server at rs.internic.net for Internet Information or the whois server
at nic.ddn.mil for MILNET Information.
```

So from the whois information we can see that the IP address of the external interface of the firewall is 193.125.25.10 We also see that the Internet Service Provider is Pipex. We can then use SamSpade's Dig feature to see what information we can recover from the DNS:

```
11/10/01 11:53:53 dig giac.co.uk @ ns.pipex.net
Dig giac.co.uk@ns1.pipex.net (158.43.192.7) ...
```

Last printed 3/8/05 8:01 PM   Page 45 of 65

```
Authoritative Answer
 Query for giac.co.uk type=255 class=1
  giac.co.uk MX (Mail Exchanger) Priority: 10 gate.giac.co.uk
  giac.co.uk MX (Mail Exchanger) Priority: 120 relay1.pipex.net
  giac.co.uk MX (Mail Exchanger) Priority: 130 relay2.pipex.net
  giac.co.uk MX (Mail Exchanger) Priority: 210 sun3.nsfnet-relay.ac.uk
  giac.co.uk MX (Mail Exchanger) Priority: 220 sun2.nsfnet-relay.ac.uk
  giac.co.uk MX (Mail Exchanger) Priority: 300 ben.britain.eu.net
  giac.co.uk MX (Mail Exchanger) Priority: 310 eros.britain.eu.net
  giac.co.uk NS (Nameserver) gate.giac.co.uk
  giac.co.uk NS (Nameserver) stile.giac.co.uk
  giac.co.uk NS (Nameserver) ns0.pipex.net
  giac.co.uk NS (Nameserver) ns1.pipex.net
  giac.co.uk SOA (Zone of Authority)
        Primary NS: giac.co.uk
        Responsible person: root@gate.giac.co.uk
        serial:2001051602
        refresh:28800s (8 hours)
        retry:7200s (2 hours)
        expire:864000s (10 days)
        minimum-ttl:86400s (24 hours)
  giac.co.uk NS (Nameserver) gate.giac.co.uk
  giac.co.uk NS (Nameserver) stile.giac.co.uk
  giac.co.uk NS (Nameserver) ns0.pipex.net
  giac.co.uk NS (Nameserver) ns1.pipex.net
  gate.giac.co.uk A (Address) 193.125.25.10
  relay1.pipex.net A (Address) 158.43.128.81
  relay2.pipex.net A (Address) 158.43.128.81
  stile.giac.co.uk A (Address) 193.125.25.11
  ns0.pipex.net A (Address) 158.43.128.8
  ns1.pipex.net A (Address) 158.43.192.7
```

From this listing we can see that as well as the gate firewall we already knew about there
is also stile.giac.co.uk at 193.125.25.11. This could give an indication that it is worth
scanning the entire 193.125.25.0/24 subnet but we have been asked to confine our
activities to the primary firewall.

3.2.1.1.1 External Port Scan[xii]

We can next run a port scan against the external interface of the firewall to see what
ports are listening. Of course it may be the case that the border router is filtering our
probes so it may be that the firewall is more vulnerable than would appear from this
external scan (in fact, the border router will allow all TCP ports and protocols through
providing they are from a valid IP address – it is merely packet filtering).

Our tool of choice for port scans is nmap. This gives us many options to scan ports

```
# nmap (V. 2.54BETA26) scan initiated Sat Nov 10 09:33:39 2001 as: nmap -sS -O -F -v -v -
P0 -oN gate.giac.co.uk.log -oG gate.giac.co.uk.grep gate.giac.co.uk
Warning:  OS detection will be MUCH less reliable because we did not find at least 1 open
and 1 closed TCP port
Interesting ports on gate.giac.co.uk (193.125.25.10):
(The 568 ports scanned but not shown below are in state: filtered)

Port        State         Service
21/tcp      open          ftp
23/tcp      open          telnet
25/tcp      open          smtp
80/tcp      open          http
1024/tcp    closed        kdm
1025/tcp    closed        listen
1026/tcp    closed        nterm
1030/tcp    closed        iad1
1031/tcp    closed        iad2
1032/tcp    closed        iad3
1058/tcp    closed        nim
1059/tcp    closed        nimreg
1067/tcp    closed        instl_boots
1068/tcp    closed        instl_bootc
1080/tcp    closed        socks
1083/tcp    closed        ansoft-lm-1
```

```
1084/tcp   closed      ansoft-lm-2
1103/tcp   closed      xaudio
1109/tcp   closed      kpop
1110/tcp   closed      nfsd-status
1112/tcp   closed      msql
1127/tcp   closed      supfiledbg
… and so on for another few hundred ports…

Too many fingerprints match this host for me to give an accurate OS guess

# Nmap run completed at Sat Nov 10 09:37:28 2001 -- 1 IP address (1 host up) scanned in
229 seconds
```

In this case we see that our expected 25/tcp (smtp) and 80/tcp (http) ports are open, but
also 21/tcp (ftp) and 23/tcp (telnet). We can confirm the default protocols are running on
these ports by telnetting to them.

We might also run vulnerability scanners such as NetRecon[xiii], SAINT[xiv] or Nessus[xv].

3.2.2 Architectural Review

This is very much a paper exercise with our auditors comparing our firewall design and
implementation against industry best practice. They should consider the following
questions:-

- **Hardware Platform:** Is our choice of platform appropriate? Would we have
  been better to use an NT4 server, given the available skill set? Would an
  appliance based firewall have been more appropriate? Does the platform offer
  sufficient resiliency and scalability? What is the total cost of ownership?

  Conclusion: A Windows 2000 platform is not the most secure basis for a
  firewall, although we appreciate the familiarity argument. Experience will be
  gained with an appliance-based firewall because one has been implemented to
  separate the internal and service networks and in the long term this technology
  provides a compelling upgrade path. The 1.2 GHz PC on which the software is
  running provides sufficient CPU to support a T1 capacity link in the near term, in
  fact it could probably cope with VPN traffic as well although this is currently
  handled by another host.

- **Firewall Software:** Have we chosen a suitable firewall package? Does
  Symantec Enterprise Firewall have an appropriate feature set? In a market where
  major players are withdrawing their firewall offerings (e.g. NAI Gauntlet) does
  Raptor have a reasonable life expectancy?

  Conclusion: Symantec Enterprise Firewall has a reasonable track record, but is
  certainly not as well established as Checkpoint Firewall 1 or Nokia Pix.
  Although it claims to integrate well with Intruder Alert IDS, our opinion is that
  Intruder Alert is a very rudimentary product and should not drive the choice of
  firewall software. Our particular dislike of Intruder Alert is that rules are stand-
  alone and there is no capability to combine rules such that a combination of
  events can trigger an alert.

- **Firewall Implementation:** How much of a risk is the single firewall (single
  point of failure?) Would it be preferable to adopt a nested design rather than run
  four subnets from one firewall host?

Conclusion: The firewall host is a single point of failure and controls several different workflows; remote staff access to the network, on-site staff access to the web and email, public access to the web server and partner VPN access to the web server. It also has scope to be misconfigured because several interfaces are controlled in one rule-set. Separate, nested, firewalls would provide a logically simpler configuration but at a higher cost that may not be justified given the relatively low traffic levels.

- **Firewall Configuration:** Considered in a subsequent section – have we configured it correctly?

### 3.2.3 Internal Network Security Review

The Internal Network Security Review will adopted many of the same principles as the External review except that the auditors are working with knowledge of, and access to, our internal network.

Network mapping and firewall interface port scans are performed for each subnet, trying to access the other subnets. We checked to see which ports are open, which services are running on them and whether they can be subverted. As before, we used tools such as nmap and nessus to perform these vulnerability tests.

Now we know that the firewall is also acting as the internal and external DNS we could interrogate it from inside and outside the firewall to see what information is available about GIAC Enterprises' hosts. We used SamSpade for this, giving the nameserver address as that of the firewall. From the external interface we got back our public server addresses whereas from the internal interface we retrieved all of the addresses we set up in section 2.5.8. We were able to confirm that there are no internal addresses being exposed in the external list.

If the scope of the audit were wider we would investigate which hosts can be reached and influenced from the internal and DMZ subnets. For example, could a disgruntled employee subvert the mail server to send abusive mail?

3.2.4 Audit of Firewall Configuration
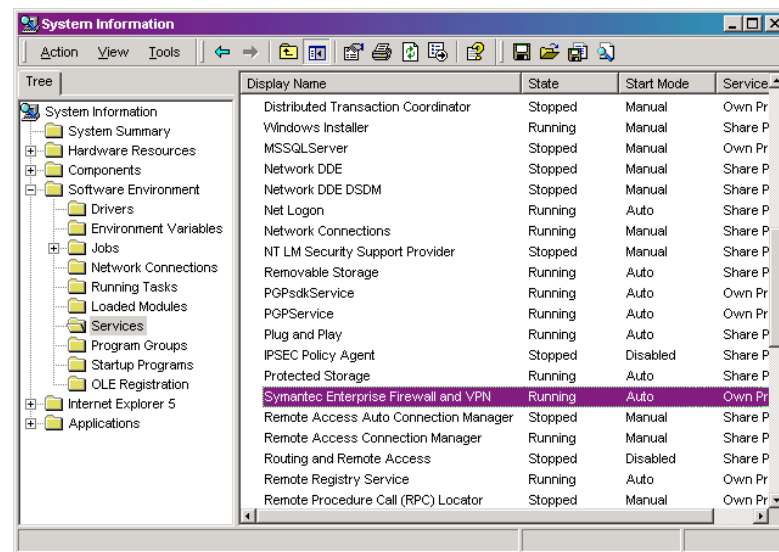
We used the WinMSD utility to report the configuration of the firewall server. This application includes the ability to produce a comprehensive configuration report, which we printed out, signed and dated and taped to the inside of the system case.

Given the list of rules we validated each rule by sending the expected traffic to the appropriate host to confirm that it works as expected. We tested this both with the native applications (e.g. mail client to SMTP server) and by using telnet to connect through the firewall to the appropriate server port and confirm the traffic reaches its intended destination. Our second laptop, running netmon, can acted as a sniffer to capture and examine packets entering and leaving the firewall.

In the event that any traffic was not correctly routed we could bug-fix our rules by enabling and disabling individual rules as well as checking their order to confirm which other rule is preventing the rule in question from working as planned.

Our audit reveals that the firewall is correctly configured, permitted traffic reaches its destination and non-permitted traffic is dropped and logged.

3.2.5 Process Audit

We audit the management practices and processes around firewall administration by checking for the existence and training in standard operating procedures, reviewing change logs and shadowing administrators as they perform their daily tasks.

This is an area in which GIAC Enterprises proves to be weak. Because only one person is responsible for IT they hold a lot of information in their head rather than formally recording it. Processes are very casual *"HR give me a ring when they take on someone new and I set up an account"* similarly the administrator can't see why she needs to record configuration changes *"No-one else here would understand what I've done in any case"*.

3.3 Evaluate The Audit

3.3.1 Risk Analysis

Our risk analysis is presented here in tabular form. It was also presented to senior management as a written report and powerpoint presentation along with the recommendations for improvement covered in the next section.

The table presented here is very much a high level summary and an in-depth discussion of risk analysis is beyond the scope of this assignment. Interested readers are recommended to read *Information Security Risk Analysis* by Thomas R Peltier[xvi]

Threat	Vulnerability	Impact	Cost	Mitigation	Recommendation
Hardware failure	Firewall is single point of failure	Unable to do electronic businesss	£30,000 per day	Alternative firewall configurations	Consider redundant firewall appliances
Internet connection failure	Discounting VPN, GIAC has single internet connection.	Staff unable to browse Internet.	Minimal	Re-route email and critical traffic through VPN firewall link, may dictate VPN being off line for part of day (afternoon)	Make contingency and disaster recovery plan to allow Internet connections to be reconfigured – and test them out of hours.
		Unable to send/receive email	£10,000 per day in lost orders and delayed invoices		
Unauthorised access to firewall and internal hosts.	Telnet and FTP ports are open on external interface.	Route of attack by hackers, potential corruption of data and systems	Depends on nature of interference with data and systems.	Disable these unused ports.	Disable these unused ports.
Hackers exploit operating system vulnerabilities	Windows 2000 runs a vast range of services by default, compared to hardened operating systems used by vendors of firewall appliances.	Much resource must be invested in tracking, understanding and protecting against new vulnerabilities.	One full-time equivalent member of staff - £100,000 per year.	Install a firewall appliance on which there is less to configure.	Replace Windows 2000 firewall with appliance based firewall.
Symantec withdraw support for Raptor	Raptor does not have a large market share and Symantec has some redundancy between its own and Axent product lines.	Firewall vulnerabilities will not be fixed forcing emergency switch to alternative product line with risk of introducing vulnerabilities	£10,000 to switch to an appliance, or unknown cost depending on at what point support is withdrawn	Either swich now to alternative platform or at least make contingency plans to do so. Enterprise-class support contract.	Plan to move to firewall appliance sooner rather than later.

Firewall misconfiguration allows unauthorised access	One firewall managing access to/from four subnets is prone to error	Variable – from no impact if misconfiguration is not exploited to complete loss of service	£0 to £30,000 per day. Assume 10% chance of causing 2 day outage in the course of a year = £600	Strict change control enabling misconfiguration to be rolled back.	Consider redundant firewall appliances

Table 1: Risk Analysis of GIAC Enterprises' Primary Firewall Architecture

3.3.2 Recommendations For Improvement

Our recommendations are in three parts

- Immediate remediation of configuration errors
- Disaster recovery and contingency planning
- Migration plan for firewall to resilient architecture

3.3.2.1   Remediation of Configuration Errors

During the audit of the firewall we found that the Telnet and FTP ports were open on the external interface. These services should be disabled as follows:

In Control Panel select Administrative Tools then Services. This displays the dialog on the left showing the services running on the computer. Navigate to the Telnet service and double click it.
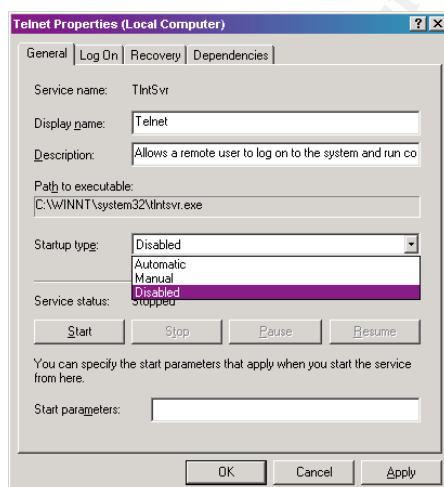
Change the start up type to disabled.

Repeat for the FTP service.

This work has zero cost or resource requirements so should be undertaken immediately.

Last printed 3/8/05 8:01 PM   Page 51 of 65

### 3.3.2.2   Disaster recovery and contingency planning

This is a question of education and awareness. The only IT resource in the company is largely self-taught and is now responsible for systems that are business critical. Should she be unavailable the systems are largely undocumented and unsupported. Consequently we recommend that additional contract resource is retained to firstly document the systems and establish a change control process (i.e. a technical author) and secondly to provide second line support on a consultancy basis (i.e. an IT security engineer/consultant)

3.3.2.3   Migration plan for firewall to resilient architecture

In our opinion, the greatest weakness of the design is its dependence on a single Windows 2000 based firewall. Host-based, rather than appliance-based, firewalls are problematic because generic operating systems tend to install and enable all kinds of services that are not required to run the firewall. If these systems are properly configured then they can be at least as secure as pre-configured appliances, but that is entirely dependent on how well they are configured.

Secondly, the single firewall introduces a single point of failure that can result in system unavailability. Whether this is acceptable depends on for how long GIAC Enterprises is prepared to be off-line. A worst case scenario could involve sourcing replacement hardware and rebuilding the firewall. In this case a commodity PC may be an advantage over an appliance because it is easier to go out and buy a replacement PC "off the shelf". On the other hand larger Value Added Resellers (VARs) should have appliances in stock.

An alternative approach is to use the configuration depicted on the left and install two firewall appliances with the service network between them. This would allow the routers to use either path to the service network providing redundant routes in the event of firewall failure.

Note that this design does not introduce additional security because compromise of either firewall will give access to the private subnets.

The single firewall also introduces a single point of failure that can result in hackers having free access to any of our private systems (other than those on the secure subnet).
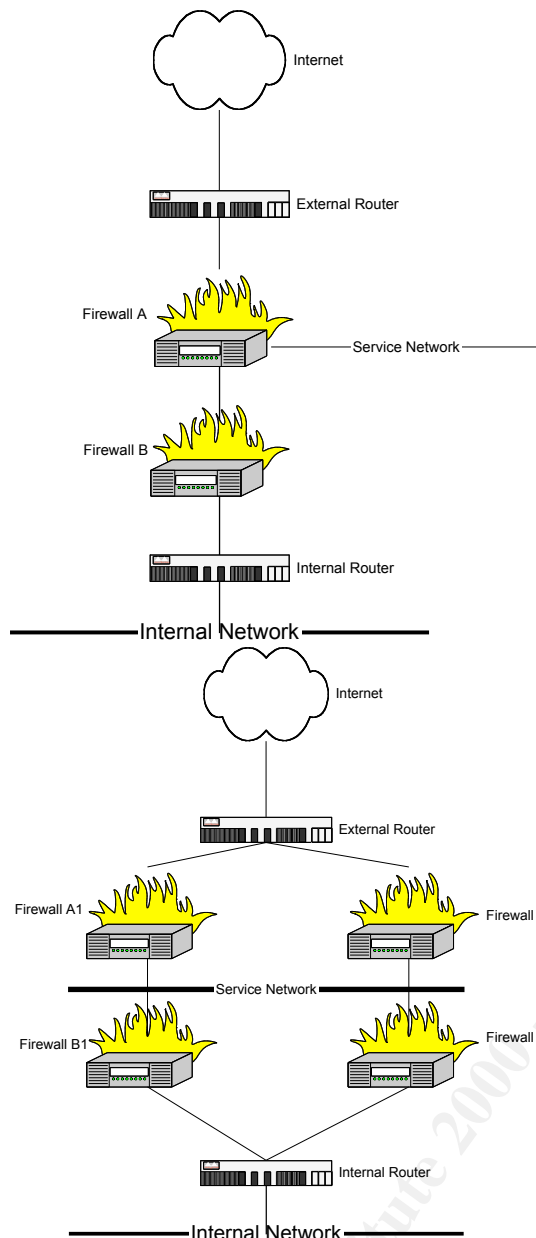
We can mitigate this risk by configuring firewalls in series so that both must be compromised before the internal network is reached. If different firewalls are used for the two layers then no single vulnerability will result in security being breached. However, either firewall is now a single point of failure that would result in the system being offline.

This final design is both resilient and layered so provides redundancy and security. However this comes at a price – we have now quadrupled the number of firewall appliances and at least doubled their management effort (assuming we use different products for the two layers). This is almost certainly overkill for a company the size of GIAC Enterprises.

Given the financial and resource constraints likely to be imposed on us by GIAC Enterprises' management, our recommendation would be to consider replacing the Windows 2000/Raptor firewall with a Nokia IP3330/Checkpoint firewall in the same configuration. The list price of such a device is $5000[xvii] which would be justified by the administrative resource saved by using a dedicated appliance. This does not address single points of failure but this could be partially mitigated by agreeing a service and maintenance contract with a Nokia VAR such that the appliance would be replaced within an agreed, short, period.

We would also encourage the IT administrator to monitor security resources such as Cert, SANS, Bugtraq and Phoneboy in order to have early warning of any vulnerabilities reported for the firewall software or platform.

4   Assignment 4 – Design Under Fire

The purpose of this exercise is to help you think about threats to your network and therefore develop a more robust design. Keep in mind that the next certification group will be attacking your architecture!

Select a network design from any previously posted GCFW practical (http://www.sans.org/giactc/gcfw.htm) and paste the graphic into your submission. Be certain to list the URL of the practical you are using. Design the following three attacks against the architecture:

1. An attack against the firewall itself. Research and describe at least **three** vulnerabilities that have been found for the type of firewall chosen for the design. Choose **one** of the vulnerabilities, design an attack based on the vulnerability, and explain the results of running that attack against the firewall.

2. A denial of service attack. Subject the design to a theoretical attack from 50 compromised cable modem/DSL systems using TCP SYN, UDP, or ICMP floods. Describe the countermeasures that can be put into place to mitigate the attack that you chose.

3. An attack plan to compromise an internal system through the perimeter system. Select a target, explain your reasons for choosing that target, and describe the process to compromise the target.

In designing your attacks, keep the following in mind:

- The attack should be **realistic.** The purpose of this exercise is for the student to clearly demonstrate that they understand that firewall and perimeter systems are not magic "silver bullets" immune to all attacks.

- The attack should be **reasonable.** The firewall does not necessarily have to be impenetrable (perfectly configured with all of the up-to-the-minute patches installed). However, you should not assume that it is an unpatched, out-of-the-box firewall installed on an unpatched out-of-the-box OS. (Remember, you designed GIAC Enterprises' firewall; would you install a system like that?)

- You **must** supply documentation (e.g., a URL to the security bulletin, bugtraq archive, or exploit code used) for any vulnerability you use in your attack.

- The attack does not necessarily have to succeed (though a successful attack is often the more interesting approach). If, given the perimeter and network configuration you have described above, the attack would fail, you can describe this result as well.

4.1 Introduction

We chose the design recently submitted by Dennis Picket, an attendee at the SANS Baltimore 2001 conference[xviii]. This design is broadly similar to our design, with a service network, internal network and high security network. It has, however, been implemented on a higher budget with redundant border routers and service network firewalls. Consequently it gives a potential upgrade path for the network we have

recommended for GIAC Enterprises. By mounting an attack against it we are forced to consider vulnerabilities in our own present and future design. Mr Picket's network design is reproduced below:
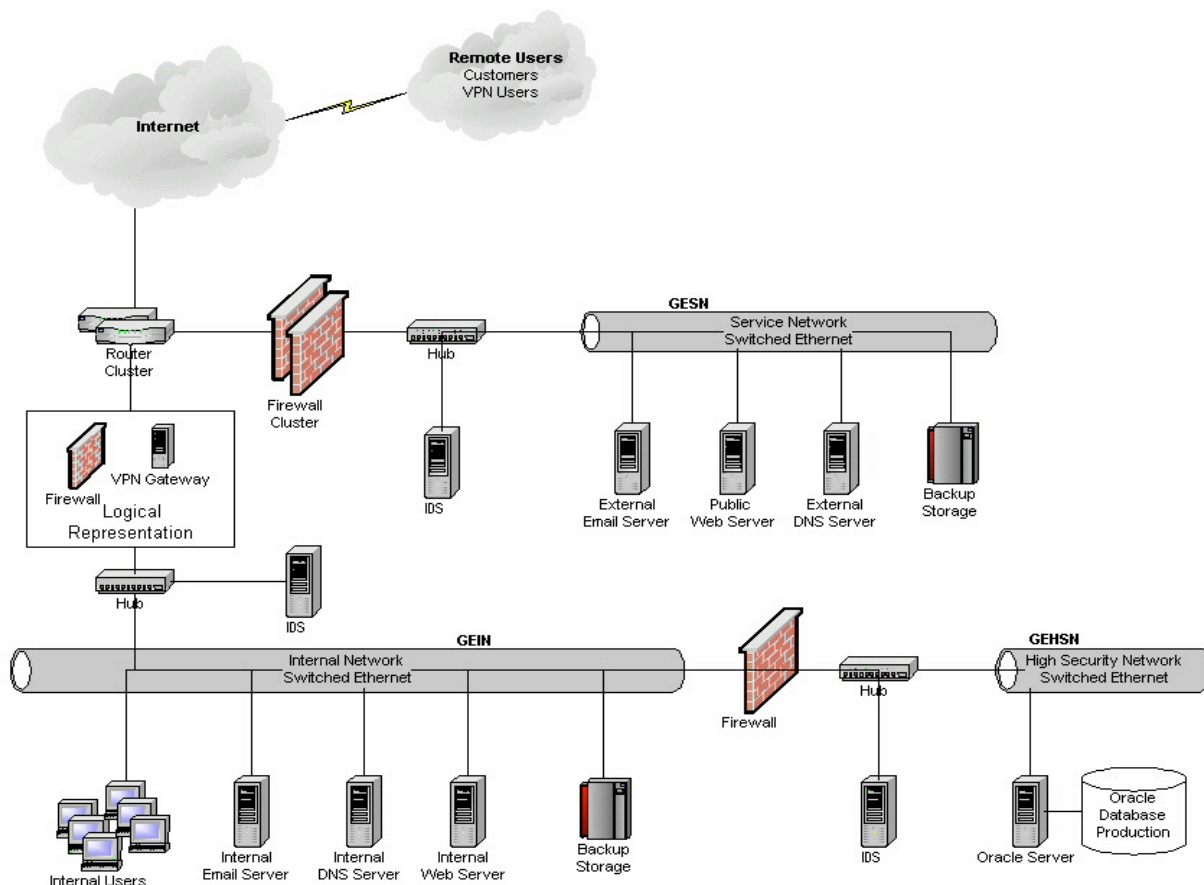
Mr Picket has specified Checkpoint Firewall 1 running on Nokia IP440s for the firewalls protecting his service and internal networks. He says

> *The Nokia IP 440 employs Check Point's **Firewall-1**[xix] firewall software, GIAC Enterprises is running the most current version, **4.0 build 4094**.*

This is somewhat surprising because his paper was submitted in the summer of 2001 and Checkpoint released version 4.1 of Firewall 1 almost a year earlier. A FAQ at Phoneboy.com relates Firewall 1 versions and builds to service packs[xx].

Build Number	Service Pack
4094	4.0 SP5
41??	4.0 SP6
4201	4.0 SP7
4304	4.0 SP8

41439	4.1 SP0
41489	4.1 SP1
41716	4.1 SP2
41814	4.1 SP3
41824	4.1 SP3 on IPSO
41862	4.1 SP4
41864	4.1 SP4 on IPSO

The table above shows all the version updates and service packs released up to the end of June 2001 – which is before Mr Picket submitted his design. Hence it would appear any vulnerabilities fixed in service packs six to eight of Version 4.0 plus any generic Version 4.0 vulnerabilities fixed in Version 4.1 could be exploited.

Information about security vulnerabilities up to v4.0 SP5 on IPSO is in a document Checkpoint released on July 26, 2000[xxi] again indicating that Mr Pickett's firewalls are a year out of date. This advisory lists the following vulnerabilities:-

- SMTP Security Server Denial of Service
- IP Fragmentation Denial of Service
- One-way Connection Enforcement Bypass
- Improper stderr Handling for RSH/REXEC
- FTP Connection Enforcement Bypass
- Retransmission of Encapsulated Packets
- Inter-module Communications Bypass
- OPSEC Authentication Vulnerability
- One-time (s/key) Password Authentication
- Getkey Buffer Overflow

It is important to emphasise that these are the vulnerabilities **fixed** by v4.0 SP5 so the system should not be prey to them. However this does put a stake in the ground marking the last set of fixes applied to the system. The system is likely to be vulnerable to any exploits reported after this.

4.2 Researching Vulnerabilities That May Still Be Active

4.2.1 Advice From The Vendor

A cynic might say that vendors only admit to vulnerabilities they have fixed – and only fix the most serious vulnerabilities. So alerts issued by Checkpoint between 26 July 2000 and the present date will highlight the most significant vulnerabilities. The Checkpoint alert archive[xxii] lists the following:

October 25, 2001
RDP Communication Issue
Check Point has become aware of a condition with RDP Protocol in VPN-1/
FireWall-1 4.1 and Next Generation (NG) that may affect system stability. If the

error occurs on a 4.1 module, certain management functions, such as logging and administrator communications, will halt. On NG modules, encryption key processing may be briefly interrupted. At no point is security compromised, and the firewall continues to enforce the security policy and allows appropriate traffic. No unauthorized access, information leakage or breach of security occurs. Check Point knows of no organizations that have had systems affected by this issue. However, Check Point recommends the hot fix below be immediately installed. QinetiQ SHC Research reported this issue to us.

July 11, 2001 (Updated September 13, 2001)
Format Strings Vulnerability
A security issue exists in VPN-1/FireWall-1 version 4.1 whereby a valid firewall administrator connecting from an authorized management client may send malicious data to a management station inside a control connection, possibly preventing proper operation of the management station. This issue exists because some instances of improper string formatting occur in VPN-1/FireWall-1 version 4.1. By sending specially constructed commands through authorized communication channels, arbitrary code may be inserted onto the operating system stack of a VPN-1/FireWall-1 management station. This vulnerability may only be exploited by an authorized and authenticated VPN-1/FireWall-1 administrator connecting from a workstation explicitly trusted by the management station, although read/write permission is not required in order to perform this attack. Since full access (read/write) administrators and those at the local system console already have direct access to the firewall system, this is an escalation of privilege only for read-only administrators.

July 9, 2001 (Updated September 13, 2001)
RDP Communication Vulnerability
Check Point uses a proprietary protocol called RDP (UDP/259) for some internal communication between software components (this is not the same RDP as IP protocol 27). By default, VPN-1/FireWall-1 allows RDP packets to traverse firewall gateways in order to simplify encryption setup. Under some conditions, packets with RDP headers could be constructed which would be allowed across a VPN-1/FireWall-1 gateway without being explicitly allowed by the rule base.

December 18 , 2000
Fast Mode Vulnerability
Check Point Software Technologies has been made aware of a TCP-fragment-based security issue associated with the use of the "Fast Mode" option for individual TCP services (NOTE: Fast Mode is synonymous with "FASTPATH" in the product GUI). … If Fast Mode has been enabled in any rule, the following issue applies. If an attacker knows the address of a protected host, or can discover it, unauthorized connection attempts can be made to that host by using a series of specially malformed TCP packet-fragments.

4.2.2 A Stateful Inspection Of Firewall 1

The phoneboy.com site also gave a link to a Firewall 1 presentation at the BlackHat 2000 conference entitled "A Stateful Inspection Of Firewall 1" by Thomas Lopatic, John

McDonald and Dug Song[xxiii]. This presentation is of interest because it uses a firewall that is identical to Mr Pickett's; a Nokia IP440 running Firewall 1 v4.0 SP5.

Using this simplified test network the authors were able to dissect the behavior of Firewall 1 in exquisite detail. Granted, the firewall is not protected from the hostile network by a border router but this would be the case if a disgruntled employee or other trusted insider makes an attack from within the Giac Enterprises network. The vulnerabilities listed in the paper are serious and summarized at length below:

Author's Note: I appreciate that in reproducing so much of this article I am standing on the shoulders of giants. All I would say is that an attacker would also be delighted to find, and use, such a clear exposition of the vulnerabilities in Firewall 1 v4.0

4.2.2.1   Authentication Attacks

In Firewall 1 v4.0 TCP port 256 is used for inter-module authentication and also by SecuRemote authentication. Thus, it is likely to be open to external as well as internal clients.

If the firewall is locally administered and has been misconfigured with

```
127.0.0.1: */none
```

in the control.map file then an attacker can spoof 127.0.0.1 and bypass authentication completely.

The verification protocol is not synchronous and the management module does not have to send its IP address before the filter module provides IP address information. Thus a bogus management module can learn the IP addresses of firewalls and genuine management modules

We can also learn significant IP addresses, such as the management module, by scanning through the address range and seeing which elicit an authentication response. The filter module does not respond to any non-significant addresses.

Firewall 1 v4.0 uses S/Key authentication as a fall back option. The implementation is flawed because it generates a new shared secret every 99 iterations and the secret is based on the time, to the nearest second. This limits the seed to 24*60*60 permutations in a day and, if the filter is polled every 10 seconds, then the 99 iterations required to force a new secret is guaranteed to happen in 990 seconds or less. Lopatec *et al*. wrote a brute force routine that was able to try all possible secrets for a given day in less than half an hour. Thus the S/Key authentication shared secret can be easily compromised.

Another authentication mechanism, FWN1, uses a secret key shared between the filter module and the management console. With FWN1 the filter generates a random number and signs it with the secret key, sending both the original number and signature to the management console. The management console is supposed to send back a different number as plaintext and signed with the same key. However no check is made that the number is different so a fake management console can just return the values it was sent in a simple replay attack.

4.2.2.2   Packet Filtering Attacks

4.2.2.2.1 TCP Fastmode

Firewall 1 had a feature called "fastmode" in which administrators could specify source or destination ports as being performance critical. Packets sent from/to these ports were then passed without being tested against the rulebase. Also, only SYN packets from fastmode ports were verified so tools such as nmap could be used to map networks through the firewall if they use a fastmode port and ensure the SYN bit is not set.

4.2.2.2.2 FWZ Encapsulation

Lopatec *et al* were also able to demonstrate an exploit of the FWZ encapsulation protocol used by SecuRemote VPN clients as follows: They found that SecuRemote packets are encapsulated by replacing the original IP destination address and protocol with that of the firewall's external interface and IP protocol 94. The trailer is then encrypted with a hash based on the IP ID.

Original Packet

Destination Address	Protocol	PAYLOAD

Encapsulated Packet

Firewall Address	IP 94	PAYLOAD	Destination Address	Protocol

Although they could not crack the key used for the hash, by holding the IP ID static they could generate a known trailer for each destination address and protocol in which they were interested and use this to craft packets. Hence they could send packets through the firewall to normally unroutable destinations such as private address ranges.

4.2.2.2.3 IP Spoofing Protection

According to Lopatec *et al* Firewall 1 versions prior to 4.1 SP1 do not implement spoofing protection against packets appearing to originate from the external interface of the firewall. As the default rule is to allow ISAKMP packets it is possible to send any UDP datagram to the external firewall interface.

They are also vulnerable to spoofed packets appearing to originate from the all-hosts multicast address (224.0.0.1). A FWZ encapsulated packet sent to this address could trick the firewall into responding to the attack host and thus initiating a connection.

4.2.3 Phoneboy Web Site

The Web site at www.phoneboy.com is a very useful resource for information about Checkpoint Firewall-1. As well as extensive lists of FAQs it also lists security alerts that have not necessarily been acknowledged by Checkpoint[xxiv] The following alerts are current at present:

> *NOTICE:* FireWall-1 4.1 SP5 (and earlier SPs) on IPSO has a problem with SYNDefender in Active Gateway mode with NAT that causes packets with untranslated addresses to leak out. A hotfix for 4.1 SP5 is available on Check Point's Software Subscription page.
>
> *NOTICE:* All versions of FireWall-1 (up to version 4.1 SP4) allow the service RDP (UDP Port 259) through the firewall by default. A hotfix is available from here. More information.
>
> *NOTICE:* If you're *not* running FireWall-1 4.0 SP7 (Solaris, NT, AIX, HPUX, Linux), FireWall-1 4.0 SP5 build 13 (IPSO), or FireWall-1 4.1 SP2 (all platforms) or later, you are vulnerable to a number of security issues. These issues were revealed at the Black Hat 2000 conference and are extremely serious in nature.
>
> *NOTICE:* A vulnerability in FAST MODE was found to exist, which people could use to get around the security policy. Note that this is not the default behavior, so you should only be vulnerable if you've explicitly enabled this feature for a TCP service. Either disable FAST MODE, upgrade to 4.1 SP3 (now available) or upgrade to 4.0 SP8 (available for all platforms except Nokia). Note that Check Point will remove this feature in the next major release since recent performance enhancements have reduced the effectiveness of this feature.

Our research has uncovered many potential vulnerabilities that we may try to exploit to attack the firewall. We need to be clear about the aims of our three types of attack; the present attack is to exploit a vulnerability in the firewall itself, the DDoS attack is to

impact the service and the attach against an internal system through the perimeter is to get within the secure perimeter.

4.3 Attack Against The Firewall

The issue is to decide on which vulnerabilities to try first in order to meet our aim of compromising the firewall. A methodical approach would be to build an attack tree as described by Bruce Schneier in his book *"Secrets and Lies"[xxv]* and in Dr Dobbs Journal[xxvi] This approach would make the desired goal "compromise the firewall" the root of the tree and approaches to achieve that goal are the leaves. By ascribing costs to the each leaf we can see which approach is the easiest or cheapest and use this first.

Unfortunately we don't have time to construct such a tree so are going to apply Occam's Razor to the vulnerabilities already listed. From the Checkpoint security alerts; the RDP issues are not believed to compromise security, the format strings vulnerability can only be exploited by firewall admins who already have privileged access, and the Fastmode vulnerability is not enabled by default.

Next we turn to the vulnerabilities listed by Lopatic, McDonald and Song. The authentication attacks they describe seem to have potential so we shall see if we can leverage these.

4.3.1 Checking For A Vulnerable Port

Communication between the management and filter modules of Firewall-1 takes place through port 256/tcp so we would first do a port scan to see if this port is listening. Given that the firewall is version 4.0 we expect that it will be because this is also the port used for SecuRemote.

4.3.2 Getting Authenticated

Our next challenge is to influence the firewall, which should require that we are authenticated. However, according to Lopatic *et al* we can issue an unload command without needing to be authenticated.

Our first attempt is to spoof the localhost address 127.0.0.1 because this is frequently misconfigured with an open allow rule to ease administration. Remember, in this case we don't need to establish two way communication, all we need to do is send the firewall a command so a spoofed source address is no problem. If the firewall had been misconfigured in the way described then we would bypass authentication completely.

If this doesn't work then we need to figure out the correct IP address of the management module. The filter module ignores authentication attempts from any hosts that do not have the correct IP address so we can scan for the correct address by trying many addresses in succession and seeing if we get a response. We will then know the IP address of a management console and can configure our hostile host to use this for subsequent attacks.

Our next approach is to attempt to brute force S/Key authentication. S/Key is the fall back option for Firewall-1 v4.0 installations without an encryption license. S/Key

generates a new key every 99 uses and they key is based on the time in seconds. Lopatic *et al* wrote a program[xxvii] to try all possible key combinations at a rate of about 50 a second so a whole days worth of secrets (24 x 60 x 60 = 86,400 combinations) could be tested in 28.8 minutes. Providing a new key is not generated in the course of the 28.8 minutes during which we are running the brute force attack we can guarantee brute forcing the key.

4.3.3 Doing Mischief

Once we are authenticated as a management console we can do some damage. Unloading the rules to leave the firewall wide open should do the trick. Lopatic *et al* wrote a program to do that too.

4.4 Distributed Denial Of Service Attack

The assignment gives us 50 "zombie" PCs with DSL connections to the Internet. If each connection has an upload speed of 250kbps then we have a total bandwidth of 12.5Mbps which could be enough to jam the internet link regardless of the particular attack we use.

A good starting point to understand denial of service attacks is Eric Cole's excellent overview in Chapter six of his book *Hackers Beware*[xxviii] A more detailed personal insight can be found on the Gibson Research Corporation web site[xxix] Steve Gibson documents a DDOS ICMP attack against his site in June 2001 and follows up with some well-targetted criticism of Microsoft's attitude to computer security and Windows XP in particular

Gibson was targeted by compromised IIS servers sending ICMP Ping packets. As we know the border router is filtering ICMP packets we will discount Ping Of Death and its relatives. They may congest the Internet link but will be stopped at the router. Assuming we wish to be more sophisticated than that we need to give the firewall something to process to tie up its resources. We can also amplify the effect of 50 PCs if we relay our attack via intermediaries broadcast addresses. Smurf does this with ICMP packets, which would probably be enough to choke the router but to reach the firewall we should use the UDP version of Smurf, Fraggle.

These amplified attacks aren't sophisticated, in fact they are the networking equivalent of assault with a blunt weapon – but if the weapon is big enough and swung hard enough it is difficult to protect against. The ogre with a stone club doesn't look for a chink in your armour he just flattens you!

If we have our choice of DDoS tools we will choose Tribal Flood Network 2000 (TFN2K) . TFN2K gives us a wide range of attacks including distributed Ping and Smurf floods (ICMP), UDP floods and TCP SYN floods – or random combinations. SYN floods can be particularly effective because they perform the first step of a three-way handshake to initiate connections with the firewall and then time out.

We can calculate the impact of a SYN flood from 50 compromised hosts: A TCP header is 20 bytes, lets call it 200 bits. Each PC has 250,000 bits of bandwidth per second so can send 1250 SYN packets. There are 50 PCs so we will be attempting to open 62,500 connections a second on the firewall.

Its important to remember that TCP SYN is a legitimate request, we can't block it without preventing all incoming connections to our network (including http and smtp). We can block unused ports so that SYN requests directed at ports on which we aren't running services will be ignored, but that doesn't help a lot because most attacks will be directed at well-known ports such as 25 and 80.

If we experience a SYN flood we can filter the source IP addresses at the router, or better yet have our ISP block them upstream. This will work for as long as the attacker doesn't change the source IP address, but if he is randomly spoofing that we can't effectively block by source address.

Perhaps the greatest contribution we can make to preventing distributed denial of service attacks is to prevent our systems from being compromised and attacking someone else. If everyone did that then there would be no problem.

4.5 Compromise A System Through the Perimeter

It is tempting to approach this task by finding some arcane trick that will allow us to sneak packets past the firewall to map out and subvert internal hosts. On the other hand it is important to realise that firewalls have to allow some traffic through in order to serve a useful business purpose. One of the business services that needs to keep running is email, and so long as security unaware staff are receiving email we have the opportunity to send them a virus or Trojan.

Of course we need to have some email addresses to send to. There are a number of ways we can get these but one of the easiest is to search on "@giac.co.uk" on a web search engine such as Google. We are likely to find all the official contact email addresses on the GIAC Enterprises web site, but its amazing how often staff use their work address as contact details for their clubs and societies or in postings to recreational newsgroups.

We may also find them in meeting attendee lists, in online LDAP directories or in commercial email databases. WS_Ping has a useful LDAP client:

Once we have some email addresses and know their format we can even guess likely ones by combining common firstnames and surnames.

When we have generated a list of likely email addresses we simply paste them into the Blind Carbon Copy field of our email, set the To: field to be an innocuous address, attach our Trojan payload and send it.

If we wanted to add a twist to this attack we could set up our own PC on the Internet with a fake @giac.co.uk email address and address book containing our other @giac.co.uk addresses as well as the email addresses of anyone else we wanted to irritate, infect our own PC with the Sircam virus[xxx] and have this send the emails. Not

only would GIAC Enterprises stand a good chance of being infected with the worm but they would most likely be blamed for other people's infections. The originating address appearing to be GIAC would also make it more likely attachments would be opened as some staff would mistake it for internal mail.

This is a shotgun approach in that we are targeting any, or all, of the internal clients but it emphasises that even with properly secured perimeters we are still at risk if our staff can be tricked. As the sage once said "*nothing is foolproof – fools are ingenious*".

5   References

[i] *"Mailgate Summary"* URL: http://www.mailgate.com/products/msummary.asp (20 October 2001)
[ii] "*The Twenty Most Critical Internet Security Vulnerabilities (Updated) The Experts' Consensus"*, The SANS Institute, Version 2.500 October 10, 2001, URL: http://www.sans.org/top20.htm (Accessed: 7 November 2001)
[iii] *" @stake LC3"*, @stake, URL: http://www.atstake.com/research/lc3/ (Accessed: 7 November 2001)
[iv] *"Cisco 1720 – Modular Access Router"*, Cisco Systems Inc., 28 June 2001, URL: http://www.cisco.com/univercd/cc/td/doc/pcat/1720.htm (20 October 2001)
[v] *"Internet Protocol v4 Address Space"*, Internet Assigned Numbers Authority, 12 September 2001, URL: http://www.iana.org/assignments/ipv4-address-space (20 October 2001)
[vi] *"Improving Security On Cisco Routers"*, Cisco Systems Inc., 15 August 2001, URL: http://www.cisco.com/warp/public/707/21.pdf (20 October 2001)
[vii] MAPS Dial Up User database, URL: http://mail-abuse.org/dul/intro.htm (14 November 2001)
[viii] nmap 2.54BETA30, Insecure.org, URL: www.insecure.org/nmap/index.html (12 November 2001)
[ix] Retina v4.7, eEye Security, URL: http://www.eeye.com/html/Products/Retina/index.html (12 November 2001)
[x] WS_ping ProPack 2.3, Ipswich Inc., URL http://www.ipswitch.com/Products/WS_Ping/ (12 November 2001)
[xi] "Chapter 3 Information Gathering" in *Hackers Beware,* Cole, Eric, (2001) New Riders Publishing
[xii] "*Port Knowledgebase"*, Network Ice, URL: http://www.networkice.com/advice/Exploits/Ports/default.htm (13 November 2001)
[xiii] Symantec NetRecon URL: http://enterprisesecurity.symantec.com/products/products.cfm?ProductID=46&PID=8883591&EID=0 (13 November 2001)
[xiv] Worldwide Digital Security Inc. SAINT URL: http://www.wwdsi.com/saint/ (13 November 2001)
[xv] Nessus URL: http://www.nessus.org (13 November 2001)
[xvi] "*Information Security Risk Analysis"*, Thomas R Peltier, 2001, Auerbach Press.
[xvii] *"Checkpoint Platform Guide"*, Checkpoint Inc. URL: http://www.checkpoint.com/products/security/platforms/platforms_list.html (14 November 2001)
[xviii] *"GIAC Certified Firewall Analyst Practical, Version 1.5e"*, Michael Dennis Picket, 15 August 2001, URL: http://www.sans.org/y2k/practical/Dennis_Pickett_GCFW.zip (28 October 2001)
[xix] www.checkpoint.com/products/firewall-1
[xx] "Which Build Number of FireWall-1 to Which Service Pack?", Dameon D Welsh-Abernathy, 29-Jun-2001, URL: http://www.phoneboy.com/faq/0385.html (28 October 2001)
[xxi] "Potential Security Issues in VPN-1/FireWall-1", Checkpoint Inc, 26 July 2001, URL: http://www.checkpoint.com/techsupport/alerts/list_vun.html, (28 October 2001)
[xxii] *"Alerts Archive"*, Checkpoint Inc, 25 October 2001, URL http://www.checkpoint.com/techsupport/alerts/index.html, (28 October 2001)
[xxiii] "A Stateful Inspection Of Firewall 1", Lopatic T, McDonald J, Song D, 26 July 2000, URL: http://www.dataprotect.com/bh2000/blackhat-fw1.html (28 October 2001)
[xxiv] *"Firewall-1 Security Alerts"*, Phoneboy, 14 October 2001, URL: http://www.phoneboy.com/homepage.html#Alerts (14 November 2001)
[xxv] *"Secrets and Lies – Digital Security In A Networked World"*, Bruce Schneier, 2000, Wiley Computer Publishing
[xxvi] *"Modeling security threats"*, Dr Dobbs Journal, December 1999, reprinted URL: http://www.counterpane.com/attacktrees-ddj-ft.html (14 November 2001)
[xxvii] Source code for these exploits is available from http://www.dataprotect.com/bh2000/blackhat-fw1.tar.gz
[xxviii] "*Hackers Beware"*, Eric Cole, 2001, New Riders Press
[xxix] *"The GRC Denial Of Service Pages"*, Gibson Research Corporation, 13 August 2001, URL: http://grc.com/dos/intro.htm (14 November 2001)
[xxx] W32.Sircam.Worm@mm, Details from Symantec Virus Database, 31 October 2001, URL http://securityresponse.symantec.com/avcenter/venc/data/w32.sircam.worm@mm.html (14 November 2001)