



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

GCFW Certification Practical

Version 1.5e

Christopher G Buckley

Document: Christopher_Buckley_GCFW.doc

Version: 1.0

Issue Date: 19 August 2001

Contents

Preface	4
Purpose and Audience	4
Organisation	4
Acronyms, Abbreviations, and Terms	4
1. Security Architecture Design (Assignments 1 to 3)	5
1.1 Business Requirements	5
1.2 Architecture Design (Assignment 1)	6
1.2.1 Requirements	6
1.2.1.1 Processes and Policies	7
1.2.1.2 Internal systems connectivity	7
1.2.1.3 Internal systems connectivity with merger partner	7
1.2.1.4 External connectivity with customers	8
1.2.1.5 External connectivity with partners	8
1.2.1.6 External connectivity with suppliers	8
1.2.1.7 Support	9
1.2.2 Proposed Architecture	10
1.2.2.1 Assumptions	11
1.2.2.2 Design Detail	11
1.3 Detailed Policy Design (Assignment 2)	15
1.3.1 Border Router (<code>router-ext</code>)	15
1.3.1.1 Increase protection for the router	16
1.3.1.2 Traffic oriented protection	18
1.3.1.3 Configure Logging	20
1.3.2 Partner/Supplier VPN (<code>vpn-conc</code>)	20
1.3.3 External DNS (<code>dns-ext</code>)	21
1.3.3.1 Split DNS	22
1.3.3.2 Install an up to date version	22
1.3.3.3 Ensure that only non -recursive service is provided	22
1.3.3.4 Limit zone-transfers	22
1.3.3.5 Disallow dynamic DNS updates	23
1.3.3.6 Ensure that only short replies are possible	23
1.3.3.7 Chroot the daemon	23
1.3.4 Mail Proxy - <code>smtpd</code> and <code>smtpfwdd</code> (<code>email-ext</code>)	23
1.3.5 Perimeter Firewall (<code>fw-ext</code>)	24
1.3.5.1 Firewall Management	27
1.3.5.2 Service Traffic	28
1.4 Audit (Assignment 3)	30
1.4.1 Audit Phases	30

1.4.1.1	Reconnaissance	30
1.4.1.2	Data Collection	31
1.4.1.3	External probing	31
1.4.1.4	Internal probing and verification of configuration	36
1.4.1.5	Presentation of results and recommendations	38
1.4.2	Financial Cost	39
1.4.3	Suggested Technical improvements	39
1.4.3.1	Resilience	39
1.4.3.2	Security	39
1.4.3.3	Forensics	40
2.	Assignment 4 - Design Under Fire	41
2.1	Social Engineering	41
2.2	Border Router	41
2.3	Perimeter Firewall	42
2.4	Exposed Services	43
2.4.1	SMTP mail	43
2.4.2	Web Server	43
2.4.3	DNS Server	44
2.4.4	VPN Server	44
2.5	Intrusion Detection System	44
2.6	Internal Firewall	45
	Appendix A. References	46

Preface

Purpose and Audience

This document contains the practical components of the GIAC Certified Firewall Analyst assessment.

It is intended for readers who have some understanding of network security issues and their interactions with business priorities.

Organisation

This document contains 2 sections/chapters in addition to this preface.

Section 1 presents a network perimeter defence architecture for a medium sized business

Section examines the weaknesses in the network perimeter architecture proposed by Ken Colson (http://www.sans.org/y2k/practical/ken_colson_gcfw.doc)

Acronyms, Abbreviations, and Terms

ARP	Address Resolution Protocol
DNS	Domain Name Service
DoS	Denial of Service
FW-1	Check Point Firewall -1
ICMP	Internet Control Message Protocol
IP	Internet Protocol
ITIL	IT Infrastructure Library
SNMP	Simple Network Management Protocol
SSL	Secure Sockets Layer
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
VPN	Virtual Private Network

1. Security Architecture Design (Assignments 1 to 3)

This section proposes an infrastructure security architecture for GIAC Enterprises, a start-up company expecting to earn \$200M a year by selling Fortune Cookie sayings online.

1.1 Business Requirements

The most important stage of any architectural design process is gathering the business requirements that must be satisfied by the final architecture. For GIAC Enterprises the following constitute the high level business requirements:

- To have a presence on the internet from which to sell Fortune Cookie sayings securely
- To have flexible, reliable and secure communications channels with customers, suppliers and partners
- To cost effectively minimise the risk associated with having an online presence
- To have cost effective and reliable internal systems

In order to derive a set of technical requirements from these business it is necessary to have more detail, however, in the absence of such detail reasonable assumptions should be made:

- The sale of Fortune Cookie sayings is not highly time sensitive i.e. sayings tend to be bought in bulk well in advance of the time they are actually required
- Suppliers and partners are unlikely to be willing to invest heavily in proprietary technology to facilitate communications and trading with GIAC Enterprises
- The primary network channel for communication with customers will be a web site
- The primary network channel for communication with suppliers will be a web site, however, in some cases wider system access may be required
- The primary network channel for communication with partners will need to be more flexible than a web site
- There will be relatively few partners but many suppliers (mostly home based)

- GIAC Enterprises works very closely with its partners and needs to provide access to wide range of internally held information
- The most sensitive information that is likely to be held by GIAC Enterprises will be customer details, order history, etc.
- Internal staff need limited access to the internet for business -relevant web browsing, email, etc.
- The nature of the GIAC Enterprises Fortune Cookies saying business is that the core company acts in a similar way to an exchange. It is likely, therefore, that the number of internal staff and IT systems will be low
- GIAC Enterprises wishes to maintain direct control over its IT infrastructure and, in particular, its security. However costs must be controlled and no vendor should have undue influence
- Due to financial constraints, the emphasis should be on building perimeter protection, but with sufficient flexibility to extend the solution inwards (such as installing firewalls to segment the internal network)

1.2 Architecture Design (Assignment 1)

This section details the proposed perimeter defence architecture for GIAC Enterprises and explains how it satisfies the business requirements.

1.2.1 Requirements

In order to produce a final architecture design, the business requirements detailed in Section 1.1 must be translated into technical requirements for each area of connectivity to be addressed.

Logically the areas to be addressed are:

- Processes and Policies
- Internal systems connectivity
- Internal systems connectivity with merger partner
- External connectivity with customers
- External connectivity with partners
- External connectivity with suppliers
- Support

1.2.1.1 Processes and Policies

No matter how comprehensive an architecture is, in any other than the very short term, the time, effort and money spent designing and building it will be wasted unless similar levels of effort are expended on designing the policies, processes, procedures and support model to accompany it. Ideally these should be based on recognised best practice (e.g. ITIL) and conform to BS7799 governing information security management. The design of such an operational model is, however, outside the scope of this document.

1.2.1.2 Internal systems connectivity

Typically, internal systems are subject to less stringent network security than those at the perimeter of the network. They provide the core business processing functionality and those ancillary services that are required to run a modern, internet connected, organisation. Typical functionality that would be present includes email, file and print services, databases, accounts systems, name services, etc. There are a number of factors governing the design of the connectivity of these systems:

1. Internal network security (as opposed to application level controls) is not a priority at this time
2. The design should be flexible enough to support the addition of network security at a later date
3. The design should allow network resources to be targeted appropriately

1.2.1.3 Internal systems connectivity with merger partner

During merger activity it is important that the IT infrastructure does not impede the business processes. It is likely that there will large information flows through the business and there will need to be widespread access to a wide range of internal systems in both organisations. Unfortunately, often the security architecture and processes of one of the merging organisations will not be of the same standard as those in the other organisation. This may be due to differing business priorities, etc. before the merger but must still be taken into account and compensated for as far as possible until a single approach can be developed. The following requirements arise from this:

1. Connectivity between defined systems within the merged businesses needs to be largely unrestricted
2. It should be possible to audit traffic between the two businesses
3. In extremis it should be possible to severely restrict or, indeed, prevent all traffic between the two businesses
4. Internal traffic must not, as far as possible, interfere or interact with external or customer traffic

5. Connectivity needs to be highly reliable as the businesses are integrated

1.2.1.4 External connectivity with customers

The primary interface to customers will be a web site. Increasingly, e-commerce web sites are reliant in back end databases for much of the functionality. Indeed some sites are entirely dynamically generated based on templates held in databases. There are therefore the following requirements:

1. Access to a front end web server should be provided
2. As appropriate, that access may be either clear text or encrypted (SSL)
3. The front end web server must be able to access its back end systems
4. It will be necessary to provide access for email to be delivered to internal email addresses
5. It will be necessary to provide name resolution for customer facing systems
6. Customers should have no access to internal systems

1.2.1.5 External connectivity with partners

The nature of GIAC Enterprises' relationships with its partners is such that the partners need to be provided with access to certain internal systems. The requirements are:

- Web site access is insufficient
- Connectivity must be provided to certain internal systems (database, email, file storage, etc.)
- Usage of such connectivity should be auditable and controllable
- Partners should not be compelled to invest in proprietary technology to facilitate connectivity

1.2.1.6 External connectivity with suppliers

This is largely the same as is required for customers (section 1.2.1.4). In some cases, however, it may be necessary to provide similar connectivity levels to those required for partners.

1.2.1.7 Support

The core business of GIAC Enterprises is the supply of Fortune Cookie sayings. Whilst unwilling to lose direct control over its IT infrastructure, the business is keen to limit its reliance on highly skilled, expensive, IT support engineers and not be reliant on a single vendor. This leads to the following requirements:

- Proprietary technologies should be avoided unless they offer significant advantages
- Both the solution as a whole and its constituent components should be relatively easy to administer and maintain
- Highly customised solutions are to be avoided as they usually result in a dependence on either consultants or a small number of key support staff
- Systems should be low maintenance

1.2.2 Proposed Architecture

The proposed architecture is detailed in Figure 1.

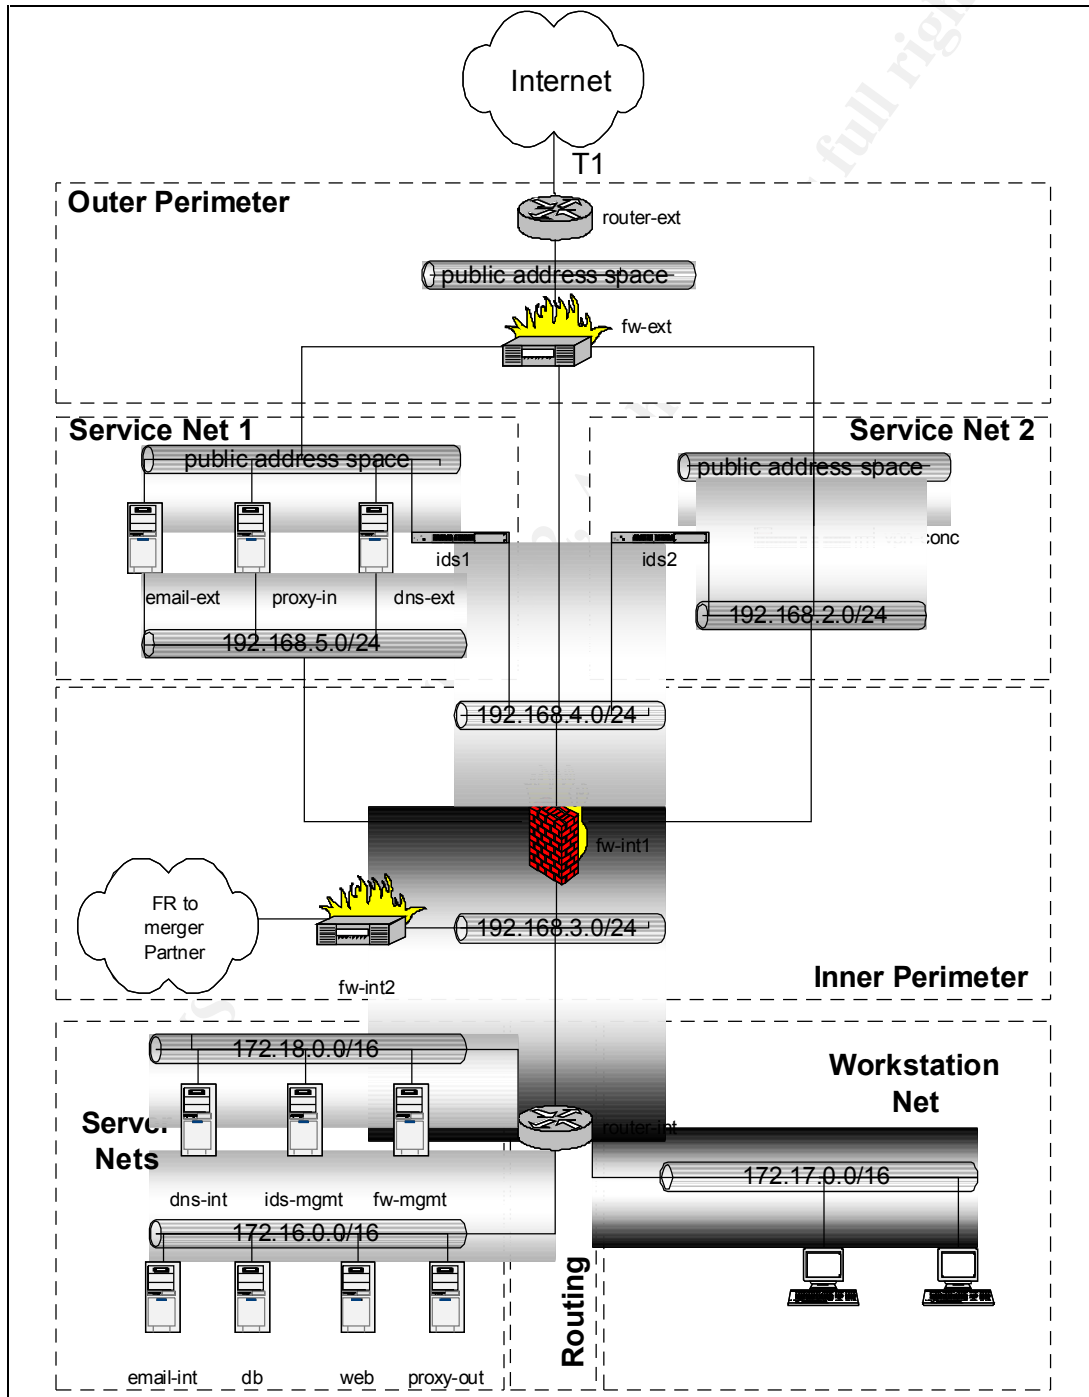


Figure 1 - Proposed Infrastructure Security Architecture

1.2.2.1 Assumptions

In producing this architecture the following assumptions have been made:

- All cross subnet network connectivity uses IP based protocols
- The two merged businesses both utilise private IP address space (Reference [9])
- The address spaces used by the merged businesses are compatible i.e. no address mapping or re-addressing is required to facilitate connectivity
- Public IP address space is available as required
- Redundancy is not required as the business is not particularly time sensitive
- The hosted applications provide their own security mechanisms (e.g. authentication, encryption, etc.)

1.2.2.2 Design Detail

This architecture adheres to a number of general principles:

- Defence in depth – always assume that someone will be able to get past the first layer of defence
- Minimisation of the number of access points to the network – monitoring of activity and analysing intrusions is easier with fewer possible intrusion points
- Functional separation – security policy enforcement (and network resource allocation) if nodes with different function types are separated within the network

The zoned nature of the design also makes it much easier to apply generalised access and administration policies. These may be based on level of risk exposure (i.e. which zone a node is in) rather than having to be written from scratch on a node by node basis.

Outer Perimeter

The outer perimeter consists of a border router (`router-ext`) and a perimeter firewall (`fw-ext`).

`router-ext` Cisco 2600 router IOS v12.2

`fw-ext` Nokia IP330 with Check Point FW -1 4.1 Service Pack 4 on IPSO

The border router provides the access point into the GIAC Enterprises Internet Channel infrastructure. The border router implements basic traffic filtering but its primary purpose is to route traffic.

The outer perimeter firewall is a stateful packet filtering appliance firewall. Its primary function is to provide filtering protection at OSI layer 3 and 4 (Network and Transport layers).

Service Nets 1 & 2

The service nets contain the nodes that will expose services to the Internet. The Service Network subnets are provided as physically separate subnets by multiple Cisco 1900 series switches.

email-ext	smtpd and smtpdwwd (from the Juniper toolkit – Reference [16]) on Sun Netra T1 (Solaris 8)
proxy-ext	iPlanet Proxy v3.6 - Sun E220R (Solaris 8)
dns-ext	ISC Bind 8.2.4 on Sun Netra T1 (Solaris 8)
vpn-conc	Cisco VPN 3030 concentrator appliance
ids1	ISS RealSecure v6.0 on Sun Netra T1 (Solaris 7)
vpn-conc	Cisco 2600 router IOS v12.2
ids2	ISS RealSecure v6.0 on Sun Netra T1 (Solaris 7)

The two service networks share common perimeter firewalls. Although the Cisco VPN device has some firewall capabilities sharing firewalls provides greater control over access to exposed nodes and eases log analysis and correlation by ensuring that data relating to particular perimeters can be found in the same logs.

Service Net 1 contains those nodes that are to be available for general use. The full details of the hardening of these nodes is out of the scope of this document. The general principles used, however, are as follows:

- Minimisation – the elimination of software and services that are not explicitly required on the host
- Management of privilege – all processes should run at the lowest privilege possible (e.g. applications running under low privilege dedicated IDs, elimination of setuid executables, etc.)
- Control over access – all administrative access to the nodes should use strong authentication and be encrypted if possible (e.g. use Secure Shell for access)

- Monitoring – each host should be actively monitored to ensure that any intrusions or unauthorised changes are spotted quickly (e.g. use of Tripwire)
- Defence in depth – each host will also have the ipfilter firewall (<http://www.ipfilter.org/>) installed and configured to ensure that it only accepts expected traffic

For further details see the following links for examples and further guidelines:

- Titan security toolkit - <http://www.fish.com/titan/>
- JASS JumpStart Architecture and Security Scripts (Toolkit) - <http://www.sun.com/blueprints/tools>
- Solaris Operating Environment Minimization for Security - <http://www.sun.com/blueprints/1299/minimization.pdf> -
- Solaris Operating Environment Network Settings for Security - <http://www.sun.com/blueprints/1299/network.pdf>
- Solaris Operating Environment Security - <http://www.sun.com/blueprints/0100/security.pdf>
- YASSP ("Yet Another Solaris Security package") - <http://www.yassp.org/>

All these nodes are dual homed and non -routing. Hence no traffic entering Service Net 1 can reach the next layer of protection without being re -initiated on one of the nodes.

Other than the IDS sensor these nodes have valid ext ernally routable public IP addresses:

- SMTP email relay (`email-ext`) - configured to allow internal email out but all externally sourced email must be destined for an internal address
- Incoming web proxy (`proxy-in`) - reverse proxy for the internal web server. Presents the web site to external browsers and terminates SSL connections
- External DNS server (`dns-ext`) - implementing split DNS, this DNS server is securely configured and only contains address entries for the Service Net hosts
- Network IDS sensor (`ids1`) – in conjunction with the firewall logs, provides warning of attacks on the Service Net 1 hosts. The sensor monitors the subnet by means of a Shomiti Tap and an interface that has no address bound to it. The sensor's second interface connects into a separate, private address space, subnet in order to connect back to its event collector

Service Net 2 is dedicated to the provision of VPN services. The VPN appliance is kept separate from the other systems due to the relatively high degree of trust that is given to connections coming from it. In the event of a node being compromised in Service Net 1, this will prevent that node being used as a platform to attack the VPN appliance directly. It also prevents users of the VPN directly accessing the other Service Net servers except via the expected routes.

- VPN appliance (`vpn-conc`) has two interfaces. The first has a public, routable IP address in order to expose its services to the Internet. The second interface is connected to an internal subnet with a private IP address. The appliance also manages a pool of available private IP addresses that it can allocate to incoming connections
- Network IDS sensor (`ids2`) – in conjunction with the firewall logs, provides warning of attacks on the Service Net 2 hosts. The sensor monitors the subnet by means of a Shomiti Tap and an interface that has no address bound to it. The sensor's second interface connects into a separate, private address space, subnet in order to connect back to its event collector

Inner Perimeter

The Inner Perimeter provides defence in depth and the ability to allow logging of all traffic that attempts to enter the internal networks.

<code>fw-int1</code>	Symantec VelociRaptor firewall appliance with Raptor v6.5 on Linux 2.2
<code>fw-int2</code>	Nokia IP330 with Check Point FW-1 4.1 Service Pack 3 on IPSO

- Service Network inner firewall (`fw-int1`) - this firewall provides protection in the event that a Service Network node is compromised. This protection is enhanced by the fact that both the type of firewall technology and the host platform are different from those in use at the outer perimeter. In addition, due to Raptor's use of proxies, it provides a high degree of application layer security to the data entering the internal networks from the Service Networks
- Firewall to merger partner (`fw-int2`) - GIAC Enterprises is connected to its merger partner by a dedicated Frame Relay connection. This provides guaranteed service levels and is not generally susceptible to Internet DoS attacks that may cripple Internet VPN solutions. This firewall provides the ability to log and/or restrict all traffic going to or from the merger partner

Routing

Routing between internal subnets is provided by a dedicated routing device.

`router-ext` Cisco 2600 router IOS v12.2

Server and Workstation Networks

These networks separate servers from users thus allowing network resources to be more easily allocated. In addition, should internal network level access controls be required then this architecture will greatly simplify their addition at a later date.

- The workstation network is purely for end user machines
- One server network is allocated to management infrastructure servers e.g. the Firewall-1 management module server (`fw-mgmt`)
- The other server network is allocated to other servers

Depending on the requirements of the corporate security policy, there may be a requirement to segregate separate divisions of the company. If this is the case then additional switched VLANs may be implemented.

All Ethernet subnets are provided as VLANs by Cisco Catalyst 2900 XL switches. Having switched fabric provides a number of advantages:

- Increased local network performance
- Reduced exposure to the threat of sniffers, etc.

1.3 Detailed Policy Design (Assignment 2)

This section provides detailed security configuration for a selection of the infrastructure components within the architecture described in Section 1.2.2. In practice, technical security policy implementation must be based around the requirements of the corporate security policy and operational model (Section 1.2.1.1). Hence, the policies described in this section are likely to need revision before actually being deployed.

1.3.1 Border Router (`router-ext`)

The border router is used as the connection point to the Internet. Its primary function is routing, but it can also be used as a first level traffic filter. In addition it should be configured to minimise its own exposure. It should be noted, however, that often the border router is under the control of the ISP rather than the destination site and hence there may be no access to the router, or control over its behaviour.

Ideally the border router should be managed directly from a serial console.

All the command examples given in this section are intended to be executed at the IOS configuration prompt. Many of them have been derived from Reference [2].

1.3.1.1 Increase protection for the router

The router should present only those services to the world that are actually required for it to fulfil its function. Those services that are presented should not enhance the ability of an attacker to compromise the network. In addition, some services may also provide a means to attack the router itself and hence access should be restricted if possible.

The following settings should be applied in global configuration mode.

1. Although it will not directly increase the security of the router, a warning banner should be created for every login prompt warning against unauthorised access to the device. The exact wording of the banner will be highly dependent on the laws within the country in which the device is sited and hence should be composed only after legal advice has been taken (% is the banner delimiter, but it may be replaced by any character not used in the banner):

```
# banner exec % banner message for EXEC processes here %  
# banner incoming % banner message for connections from  
external hosts here %  
# banner login % banner message for login attempts here  
%
```

2. Disable the finger service as it may give out user information:

```
# no service finger
```

3. Disable the small services (echo, chargen, etc.) as they are not required:

```
# no service udp-small-servers  
# no service tcp-small-servers
```

4. Disable the IP DNS-based host lookup on the router to reduce exposure to DoS attack based on forcing large numbers of lookups:

```
# no ip domain-lookup
```

5. Disable the BOOTP service as it is not required:

```
# no ip bootp server
```

6. Disable the Maintenance Operations Protocol (MOP). This could be applied to a single interface if greater flexibility is required:

```
# no mop enabled
```

7. Disable the Cisco Discovery Protocol as it is not required:

```
# no cdp run
```

8. Ensure that the Berkeley 'r' commands are disabled on the router as they use very weak authentication:

```
# no ip rsh-enable  
# no ip rcmd rcp-enable
```

9. Disable identification support thus preventing remote sites gaining identification information about a TCP port:

```
# no ip identd
```

10. Disable the web interface for the router as it presents an extra avenue of attack:

```
# no ip http server
```

11. SNMP has been the source of a range of exploits, can give away a wide range of information and can even provide control over router behaviour. Access to it should, therefore, be controlled. The following commands define an access list and read-only and read-write SNMP communities which use the access list:

```
# access-list 10 permit <management station ip>  
# snmp-server community s1t3=S3c rw 10  
# snmp-server community n0t=S3c ro 10
```

12. Command line access to the router should also be controlled. The following commands set up both a user and Privileged Mode password and associate an access-list with the router vty lines:

```
# line vty 0 4
# login
# password <password>
# enable secret <admin password>
# access-list 20 permit <management station ip>
# line vty 0 4
# access-class 20 in
```

13. If a TFTP server is being used to store router configurations, then, due to the lack of authentication in the TFTP protocol, it must be protected to prevent outside attempts to gather the configurations or even alter them. This may be achieved by using a firewall to prevent access except from a very small range of addresses

1.3.1.2 Traffic oriented protection

The following settings reduce the exposure to specific types of traffic based attacks. The aim here is not to provide complete protection, it is to provide a first level filter rather than a comprehensive firewall configuration.

The following should all be set in global configuration context.

1. Disable support for source routing. Source routing may be used to launch attacks from unexpected directions (e.g. partners with less stringent security):

```
# no ip source-route
```

2. Disable TCP selective acknowledgement (RFC 2018. Reference [11]). This reduces performance but increases protection against DoS:

```
# no ip tcp selective-ack
```

The following should all be set in the context of the Internet facing interface.

1. Turn off directed IP broadcasts thus preventing the router becoming a broadcast amplifier in a DoS attack on a third party:

```
# no ip directed-broadcast
```

2. Disable the sending of ICMP unreachable messages for connections blocked by an access list. This will make scanning for available services on the router much slower as an attacker will have to wait for connection timeouts:

```
# no ip unreachable
```

3. Disable proxy ARP support to prevent the disclosure of internal MAC addresses:

```
# no ip proxy-arp
```

4. Implement rules to reduce the possibility of obviously spoofed traffic and ICMP redirects passing the router. This may be achieved by blocking traffic from RFC 1918 (Reference [9]) private or reserved address ranges. In addition log dropped traffic including which interface it was dropped at. This assumes that the Internet facing interface is serial0:

```
# interface serial0
# ip access-group 110 in
# access-list 110 deny ip 0.0.0.0 0.255.255.255 any log -input
# access-list 110 deny ip 127.0.0.0 0.255.255.255 5 any log -input
# access-list 110 deny ip 169.254.0.0 0.0.255.255 any log -input
# access-list 110 deny ip 240.0.0.0 7.255.255.255 any log -input
# access-list 110 deny ip 248.0.0.0 7.255.255.255 any log -input
# access-list 110 deny ip 10.0.0.0 0.255.255.255 any log -input
# access-list 110 deny ip 192.168.0.0 0.0.255.255 any log -input
# access-list 110 deny ip 172.16.0.0 0.15.255.255 any log -input
# access-list 110 deny icmp any any redirect log -input
# access-list 110 permit ip any any
```

5. Implement egress filtering to reduce the chance that spoofed traffic leaves the network. Assuming that the valid public addresses for the perimeter firewall and the service network nodes are 192.10.10.0/28 and 192.10.11.0/30 then these rules would be:

```
# interface serial0
# ip access-group 120 out
# access-list 120 permit ip 192.10.10.0 0.0.0.15 any
```

```
# access-list 120 permit ip 192.10.11.0 0.0.0.3 any
# access-list 120 deny ip any any log -input
```

6. If routing table update protocols are in use e.g. Border Gateway Protocol, OSPF, etc. then authentication of updates should be implemented to ensure that updates are only accepted from valid sources. Details of the implementation are specific to the protocol in use.

7. Ensure split horizon is enabled to reduce the chance of routing loops:

```
# show ip interface serial0
```

check for the following in the output:

```
Split horizon is enabled
```

1.3.1.3 Configure Logging

Finally, any events on the router should be logged to an external server. Here they are logged to a local log buffer, a syslog server and not to the console:

```
# logging buffered 8192 debugging
# no logging console
# logging trap debugging
# logging <target ip>
```

Where <target ip> is a public address made available by the external firewall using the proxy arp capabilities of the Nokia platform. The syslog server is inside the protected network and is made available to the router via a NAT rule on the perimeter firewall.

1.3.2 Partner/Supplier VPN (vpn-conc)

The Cisco VPN concentrator needs to be configured to allow authorised third parties to connect to GIACC Enterprises.

Using the interactive command line interface (see Reference [15] for details) the following should be configured:

- First the interfaces should be configured with valid addresses and subnet masks for the subnet they are on
- The system name and details should be assigned

Using the web interface the following should be configured (the web interface is more intuitive)

- Tunneling protocols and options should be defined
 - PPTP should be enabled and encryption required
 - IPSec should be enabled

These allow the fullest compatibility with third party clients, in particular those provided by default with Windows 2000, etc.

- Address assignment should be configured so that incoming connections are given an end point address. The method chosen should be Configured Pool Address Assignment with a start address of 192.168.2.10 and an end address of 192.168.2.240 (assuming that none of these are already utilised on the subnet)
- Routing into the internal network should be configured via static routes
- The authentication method should be set to be the internal authentication server
- The following IPSec parameters should be set
 - A shared key should be set for LAN to LAN connections
 - An IPSec group should be created that defines the Security Association (SA) to be used: ESP/IKE-3DES-MD5

ESP (Encapsulating Security Payload – IP protocol 50) both authenticates and encrypts the data being transferred. AH (Authentication Header – IP Protocol 51) can not be used if there is any possibility of network address translation occurring anywhere on the route between the client and the concentrator as the header will be modified and hence authentication will fail. In addition, AH does not provide any encryption of the data.

- Strong passwords will be enforced and the longer term aim is to move to IPSec only VPNs with SecurID token based authentication. However, until the business is better established the cost cannot be justified.

1.3.3 External DNS (dns-ext)

Historically the DNS daemon BIND has proven vulnerable to a range of exploits. In addition, poorly configured installations provide an easy way for potential intruders to reconnoitre a site. In order to reduce the risk of providing external services using BIND the following steps should be taken (configuration detail from Reference [13]):

1.3.3.1 Split DNS

As discussed earlier, DNS will be implemented in a split configuration. That is, internal queries will be served by one server and external queries will be served by another. The server for external queries will hold no information at all about any hosts other than those that external clients need to know about (e.g. the external addresses of those in the service networks).

1.3.3.2 Install an up to date version

It is recommended that ISC BIND (<http://www.isc.org/products/BIND/>) be installed rather than using the Sun supplied package. This is because, typically, security issues will be patched much faster in the ISC distribution than the Sun package. In addition, as ISC BIND is available in source, it is possible to customise the installation if required (e.g. alteration of binary and configuration file location).

At this time the most recent versions of ISC BIND are BIND 8.2.4 and 9.1.3. The configuration of these version is similar. ISC recommend that a version 9.x variant be deployed but, if version 9.x is viewed as too immature, etc. then version 8.x is still acceptable. Under no circumstances should a version 4.x variant be deployed.

1.3.3.3 Ensure that only non -recursive service is provided

The daemon should be configured to only answer queries about those hosts for which it is authoritative. Under no circumstances should the daemon reach recursively to answer queries. This is to stop the daemon being directed to query another server under the attackers control that could result in DNS cache poisoning (i.e. the insertion of incorrect entries in the name server cache).

This may be achieved by inserting the following directive in the `named.conf` file (the BIND configuration file):

```
options {  
    recursion no;  
};
```

1.3.3.4 Limit zone-transfers

Limiting zone transfers prevents third parties from downloading the full table of entries for the domain.

This may be achieved by added the following directive to `named.conf`:

```
options {  
    allow-transfer { none; };
```

```
} ;
```

1.3.3.5 Disallow dynamic DNS updates

The dynamic update facility (described in RFC 2136 – Reference [14]) enables authorised updaters to add and delete resource records from a zone for which the server is authoritative. This could, all too easily, be used maliciously if the facility were generally available.

By default BIND 8 and 9 disallow dynamic updates.

1.3.3.6 Ensure that only short replies are possible

As zone transfers are not required, we can limit which protocols we have to make available as channels over which DNS queries may be made if we can ensure that no possible reply from the daemon is greater than 512 bytes. The solution - keep host names, aliases, etc. short.

1.3.3.7 Chroot the daemon

In the event that the daemon is compromised and access gained to its host, then it is important to limit what the intruder can do. This can be achieved by running the daemon as a low privilege user from a 'chroot' jail. A chroot jail presents the daemon with the minimum system facilities (libraries, configuration files, etc.) that are required to run. To the daemon, and processes spawned from it, it appears as if the root directory of the jail is actually the root directory of the whole system – thus preventing movement outside the jail.

In BIND 8 the command line options required to implement this are:

- u specifies the username to use
- g specifies the group ID to use
- t specifies the directory for the server to chroot to

Full details, including which libraries are required if BIND is not statically linked, etc. may be found in Reference [13].

1.3.4 Mail Proxy - smtpd and smtpfwdd (email-ext)

Historically, sendmail has been the target of many serious exploits. Due to its complexity it is hard to configure let alone secure. In order to provide a higher level of security for the inbound email relay a simpler alternative will be deployed.

The tool chosen is smtpd (and its pair smtpfwdd) from the Juniper toolkit (<http://www.obtuse.com/juniper/>). This pair of daemons separate the job of receiving email (smtpd) and forwarding it on to its destination (smtpfwdd).

The configuration required for to prevent relaying to destinations outside the domain is as follows:

```
# We trust everything from inside on a trusted interface
to go out
allow:UNTRUSTED:ALL:ALL

# DNS registered clients can talk to me, with mail for my
domains
allow:KNOWN:ALL:*.giac.com

# unregistered clients get dropped.
deny:UNKNOWN:ALL:ALL

# otherwise mail to nonlocal users won't get relayed.
noto:ALL:ALL:ALL:551 Sorry %H(%I), I don't allow
unauthorized relaying. Please use another SMTP host to
mail from %F to %T
```

This assumes that the protected domain is `giac.com` .

1.3.5 Perimeter Firewall (`fw-ext`)

Check Point Firewall-1 adopts a 'deny by default' policy to network traffic. That is, if the traffic has not been explicitly allowed then it is not permitted to pass. The advantage of this is that, as long as the required traffic is tightly defined, nothing unexpected should be able to pass. An alternative approach would be to allow everything through unless it were explicitly forbidden. This, however, places the onus on the administrator to know details of every possible type of dangerous traffic – it is likely that many will be overlooked.

The rule set will be constructed in two stages:

1. Definition of rules to enable management of the firewall, whilst protecting it from intruders
2. Definition of rules to allow appropriate traffic in and out of the service networks

Firstly, however basic rule set policy must be defined (**Policy** -> **Properties**):

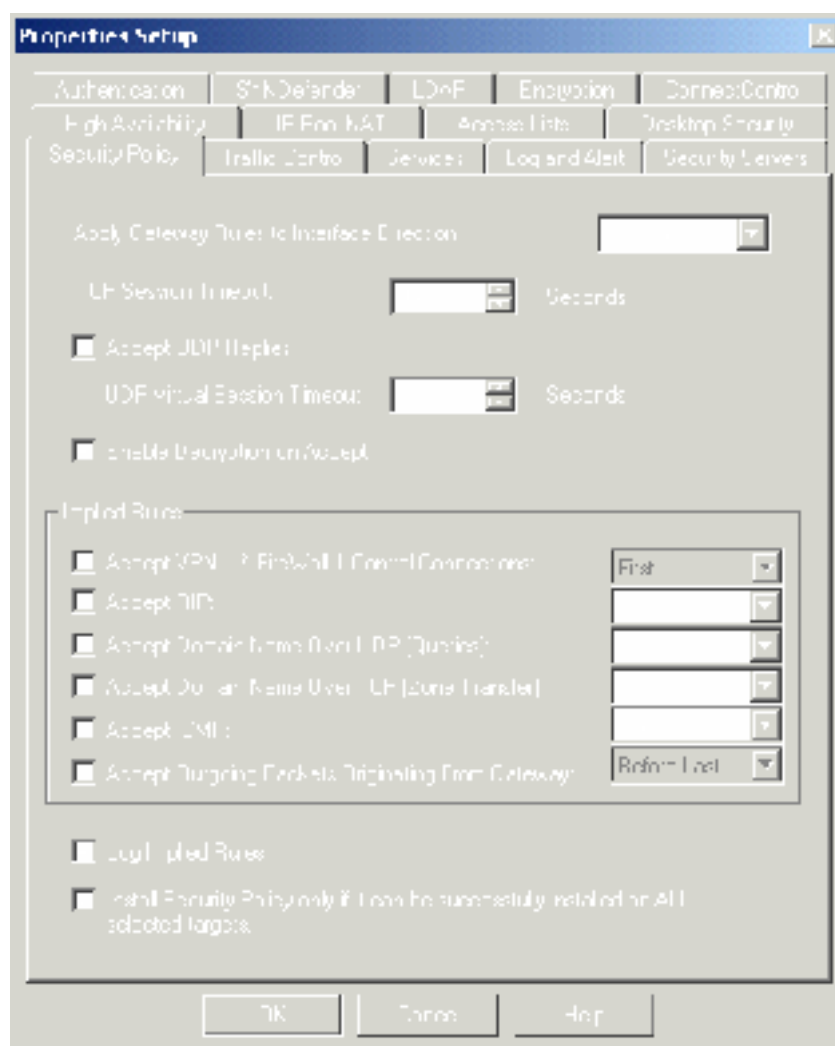


Figure 2 - Rule set properties

Next, objects must be created for each of the nodes and subnets that the external firewall is protecting.

e.g. (Manage -> Network Objects)




Figure 3 - Network Objects




Figure 4 - Define a Server

1.3.5.1 Firewall Management

The management rules are:

- The first rule ensures that all unauthorised traffic to the firewall is dropped and is logged
- Through away, un-logged, Netbios traffic from non -service network and internal addresses as it is almost akin to background noise on the Internet at this time
- Allow Firewall-1 management traffic and Secure Shell (a new service type defined as being TCP traffic to port 22) from the management module server




Figure 5 - Firewall Protection

1.3.5.2 Service Traffic

The following represents the traffic that we wish the external perimeter firewall to pass:

- Allow UDP DNS (port 53) queries to the external DNS server from any source. TCP is only used for DNS for zone transfers and large replies (in excess of 512 bytes). We don't want to allow zone transfers and are in a position to ensure that no reply will be over 512 bytes – so we can disallow TCP for DNS.
- Allow DNS queries out from Service Net addresses
- Allow DNS queries out from the internal DNS server and NAT them
- Allow SMTP (TCP port 25) traffic to the mail relay (email-ext)
- Allow SMTP (TCP port 25) traffic from the internal mail server (email-int)
- Allow HTTP (TCP port 80) connections to the proxy from external sources and NAT them
- Allow HTTPS (TCP port 443) to the proxy from external sources
- Allow outbound FTP, HTTP and HTTPS connections from the internal proxy
- Allow VPN (negotiation and IPSEC and PPTP) connections from external sources. These are:
 - Authentication Header (AH) described in RFC 2402. This is IP protocol 51
 - Encapsulation Security Payload (ESP). This is IP protocol 50
 - ISAKMP – A protocol framework for key exchange. UDP destination port 500
 - SKIP. This is IP protocol 57

- GRE (General Routing Encapsulation) IP protocol 47 (for PPTP)

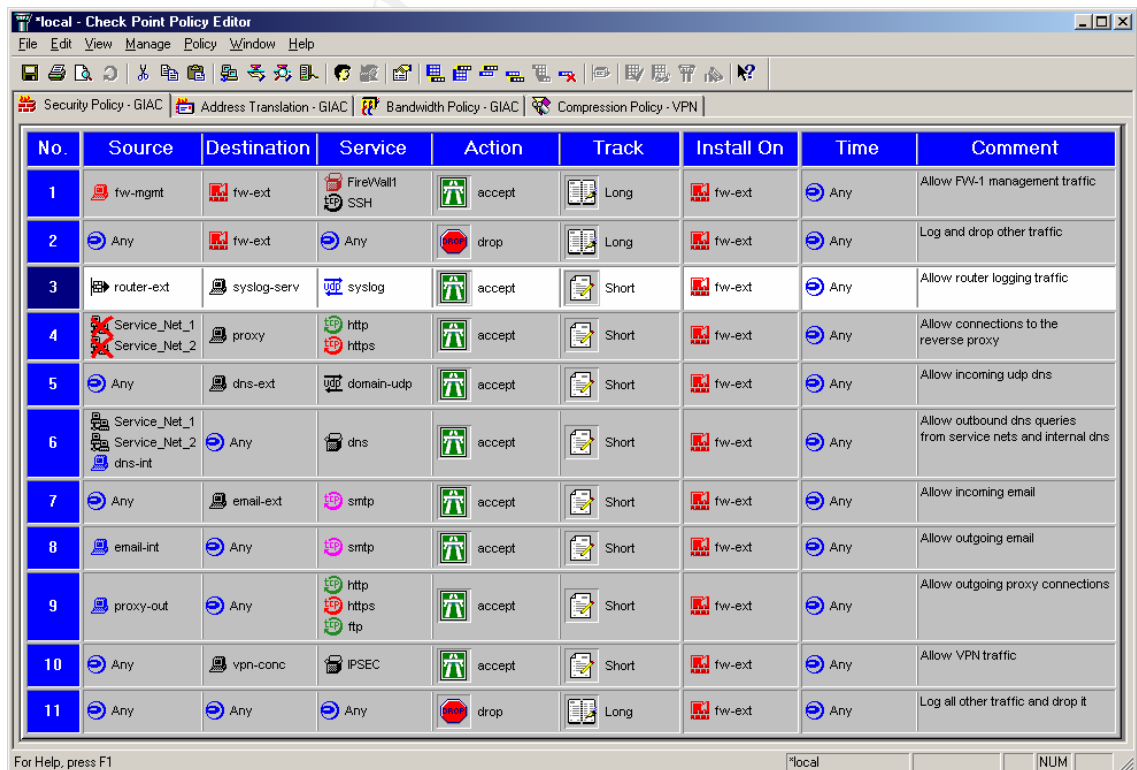
FW-1 bundles these IPSEC services into a single predefined IPSEC service object. The VPN requirements and policies will be discussed further in Section 1.3.2

- Allow syslog from the external router to the logging server via NAT
- Log all other traffic and drop it

As packets traverse the rules set from top to bottom until a match is made, the order of the rules is important from both a security and performance standpoint. Care should be taken to ensure that the intent of a rule further down the set is not masked by the action of a rule further up (though FW-1 will raise some warnings in that case). Also, where large traffic flows and long rule sets are involved, rules should be arranged such that the most important and high volume traffic matches rules early i.e. to reduce the number of rule comparisons that must be performed.

Where expected traffic is logged, logging is at lower level of detail compared to where the traffic was not expected.

The final rule set in the Firewall-1 Policy Editor, therefore, is:



No.	Source	Destination	Service	Action	Track	Install On	Time	Comment
1	fw-mgmt	fw-ext	FireWall1 SSH	accept	Long	fw-ext	Any	Allow FW-1 management traffic
2	Any	fw-ext	Any	drop	Long	fw-ext	Any	Log and drop other traffic
3	router-ext	syslog-serv	syslog	accept	Short	fw-ext	Any	Allow router logging traffic
4	Service_Net_1 Service_Net_2	proxy	http https	accept	Short	fw-ext	Any	Allow connections to the reverse proxy
5	Any	dns-ext	domain-udp	accept	Short	fw-ext	Any	Allow incoming udp dns
6	Service_Net_1 Service_Net_2 dns-int	Any	dns	accept	Short	fw-ext	Any	Allow outbound dns queries from service nets and internal dns
7	Any	email-ext	smtp	accept	Short	fw-ext	Any	Allow incoming email
8	email-int	Any	smtp	accept	Short	fw-ext	Any	Allow outgoing email
9	proxy-out	Any	http https ftp	accept	Short	fw-ext	Any	Allow outgoing proxy connections
10	Any	vpn-conc	IPSEC	accept	Short	fw-ext	Any	Allow VPN traffic
11	Any	Any	Any	drop	Long	fw-ext	Any	Log all other traffic and drop it

Figure 6 - Final rule set

No	Original Packet			Translated Packet			Install On	Comment
	Source	Destination	Service	Source	Destination	Service		
1	email-out	Any	Any	fw-out	Original	Original	fw-out	NAT outgoing email
2	dns-out	Any	Any	fw-out	Original	Original	fw-out	NAT outgoing dns queries
3	proxy-out	Any	Any	fw-out	Original	Original	fw-out	Nat outgoing proxy requests
4	router-out	syslog-srv	syslog	Original	fw-objct	Original	fw-out	Pass syslog to management server

Figure 7 - Nat Rules

1.4 Audit (Assignment 3)

Throughout the whole lifecycle of an infrastructure implementation various levels of audit are required to ensure that the solution will be both secure and manageable. This section focuses on a post implementation audit and describes a very comprehensive audit of the type that might be undertaken by regulatory technical auditors.

1.4.1 Audit Phases

The audit will involve several phases:

- Reconnaissance
- Data Collection
- External probing
- Internal probing and verification of configuration
- Presentation of results and recommendations

1.4.1.1 Reconnaissance

Once the audit has been commissioned, this is the first phase. The primary purpose is to see what information about the infrastructure can be gathered from public sources e.g.

- Mailing list archives – email headers often expose details of internal network structure
- News articles and vendor publicity – these may provide details of technologies in use within the organisation

- The company reception desk – is it possible to acquire an organisation chart or phone directory?

This phase should take 1 day.

1.4.1.2 Data Collection

During this phase data sources within the organisation should be explored such that the business aims and requirements can be better understood and the audit results placed in the correct context. Any existing security or system policies and procedures should be sourced. In addition, network diagrams, component version and patch numbers etc. should also be obtained.

It is important to remember, of course, that often documented network diagrams do not accurately reflect reality!

It is likely that, in even a relatively small organisation, this phase may well take over a week.

1.4.1.3 External probing

External probing of the network should take place using the network diagrams for reference. Although this may be thought 'unfair' it does represent the worst case scenario.

Scanning would be performed at low-impact times (e.g. outside business hours)

For generic scanning of services and vulnerabilities the following tools will be used:

- nmap by Fyodor (<http://www.insecure.org>). nmap is a port scanning and mapping tool that implements a variety of methods of avoiding perimeter protection measures e.g. SYN scans, FIN scans, etc. The following is sample output of a TCP connect scan ran against all TCP ports of an early release Solaris 8 machine with no hardening (the nmap release has since been superseded by 2.54BETA28):

```
bash-2.03# nmap -sT -O -pl-65535 192.168.2.6

Starting nmap V. 2.3BETA12 by Fyodor (fyodor@dhp.com, www.insecure.org/nmap/)
Interesting ports on (192.168.2.6):
Port      State      Protocol  Service
7         open      tcp       echo
9         open      tcp       discard
13        open      tcp       daytime
19        open      tcp       chargen
21        open      tcp       ftp
23        open      tcp       telnet
25        open      tcp       smtp
37        open      tcp       time
79        open      tcp       finger
111       open      tcp       sunrpc
512       open      tcp       exec
513       open      tcp       login
514       open      tcp       shell
```



```
515    open    tcp    printer
540    open    tcp    uucp
2049   open    tcp    nfs
4045   open    tcp    lockd
5987   open    tcp    unknown
6112   open    tcp    dtspc
7100   open    tcp    font-service
10128  open    tcp    unknown
32771  open    tcp    sometimes-rpc5
32772  open    tcp    sometimes-rpc7
32773  open    tcp    sometimes-rpc9
32774  open    tcp    sometimes-rpc11
32775  open    tcp    sometimes-rpc13
32776  open    tcp    sometimes-rpc15
32777  open    tcp    sometimes-rpc17
32778  open    tcp    sometimes-rpc19
32779  open    tcp    sometimes-rpc21
32780  open    tcp    sometimes-rpc23
32781  open    tcp    unknown
32782  open    tcp    unknown
32784  open    tcp    unknown

TCP Sequence Prediction: Class=random positive increments
                        Difficulty=38400 (Worthy challenge)
Remote operating system guess: Sun Solaris 8 early acces beta (5.8)
Beta_Refresh February 2000

Nmap run completed -- 1 IP address (1 host up) scanned in 432 seconds
```

The output clearly shows the large number of TCP services that a default Solaris 8 install makes available. These include echo and chargen (used in a classic DoS attack), login and shell (very weak authentication and trust) and a wide range of others. This machine is totally unsuited to the role of an exposed host.

Scanning of the mail host results in:

```
bash-2.03# nmap -sT -O -p1-65535 192.168.5.10

Starting nmap V. 2.3BETA12 by Fyodor (fyodor@dhp.com, www.insecure.org/nmap/)
Interesting ports on (192.168.5.10):
Port      State      Protocol  Service
25        open      tcp       smtp
22        open      tcp       secureshell

TCP Sequence Prediction: Class=random positive increments
                        Difficulty=38400 (Worthy challenge)
Remote operating system guess: Sun Solaris 8 early acces beta (5.8)
Beta_Refresh February 2000

Nmap run completed -- 1 IP address (1 host up) scanned in 440 seconds
```

Only the required services are running. The process should then be repeated but for UDP. This should be combined with running netstat on each host to ensure that the two sets of results correspond as below.

```
# netstat -an

UDP: Pw4
  Local Address    Remote Address    State
-----
  *.               Unbound
```

GCFW Certification Practical

Version 1.5e

Version 1.0

TCP: P v4							
Local Address	Remote Address	Swind	Send-Q	Rwind	Recv-Q	State	
**	**	0	0 24576	0	0	IDLE	
*.23	**	0	0 24576	0	0	LISTEN	
*.25	**	0	0 24576	0	0	LISTEN	
192.168.2.6.23	192.168.2.4.2175	16946	1 24820	0	0	ESTABLISHED	
**	**	0	0 24576	0	0	IDLE	

TCP: P v6							
Local Address	Remote Address	Swind	Send-Q	Rwind	Recv-Q	State	If
**	**	0	0 24576	0	0	IDLE	
*.23	**	0	0 24576	0	0	LISTEN	
*.25	**	0	0 24576	0	0	LISTEN	

- Nessus (<http://www.nessus.org>). Nessus is a freely available, easily extensible, vulnerability scanning tool

Nessus has a wide range of scanning plugins (see Figure 8) that test for a very broad range of vulnerabilities.

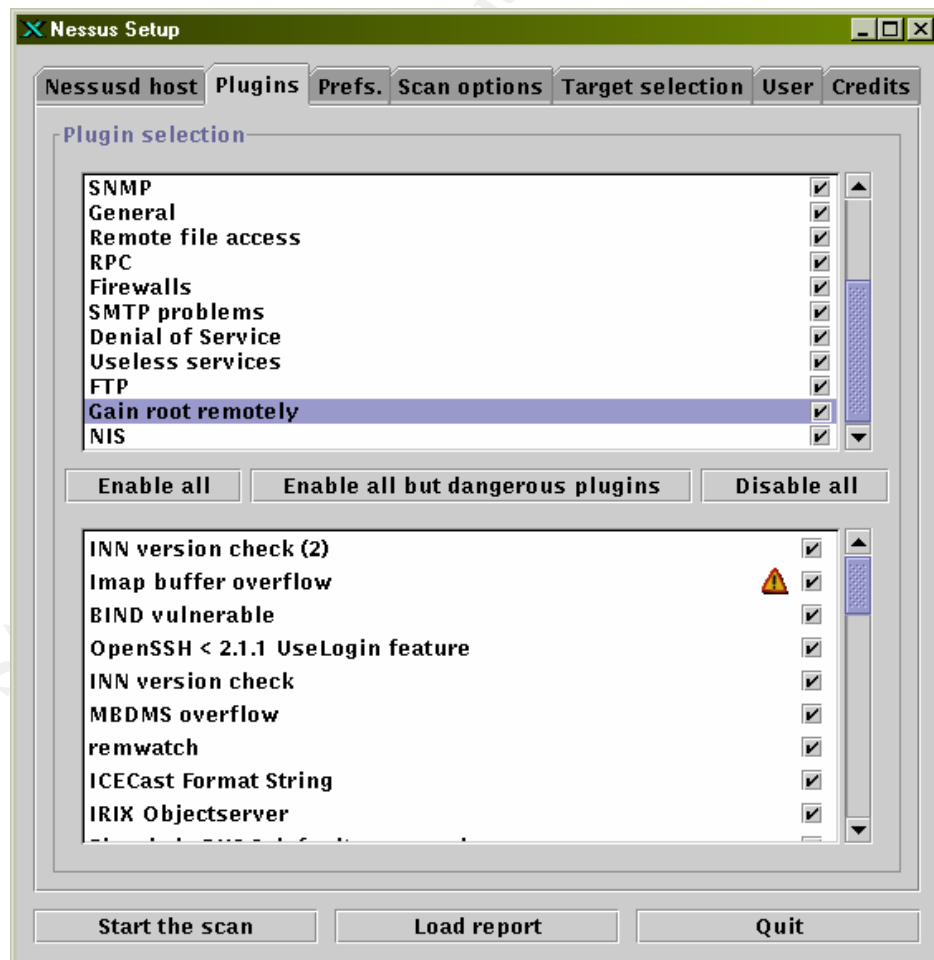


Figure 8 - Nessus Plugins

Nessus can also be configured to make use of NMAP as part of its vulnerability scan (see Figure 9)

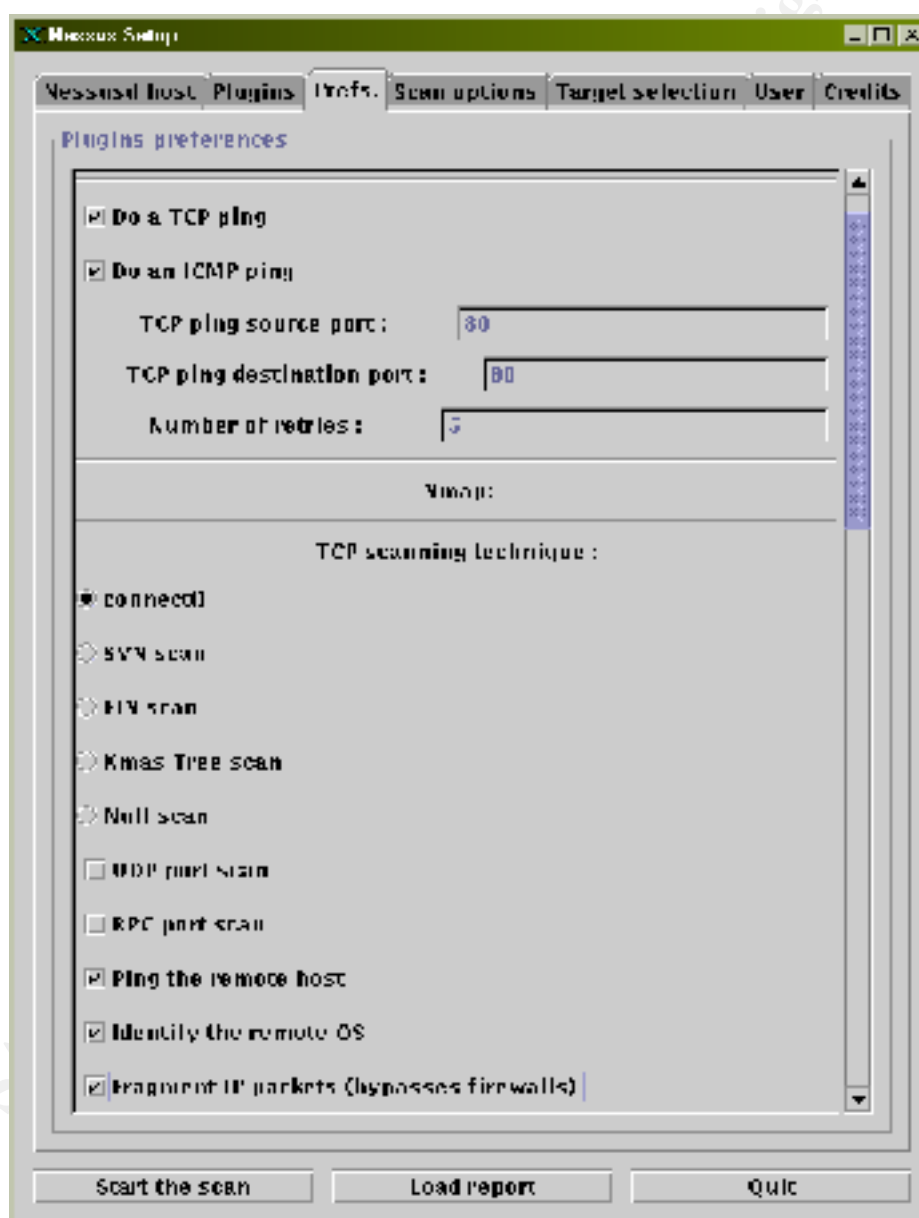


Figure 9 – Nessus NMAP usage

Nessus also has clear reporting and severity categorising of scan results. Figure 10 and Figure 11 show the results of a Nessus scan against the same host for which the NMAP scan is shown above. Care must be taken when interpreting these, however, as what may be presented as a vulnerability may also be intended behaviour permitted by the security policy (e.g. the use of finger services)

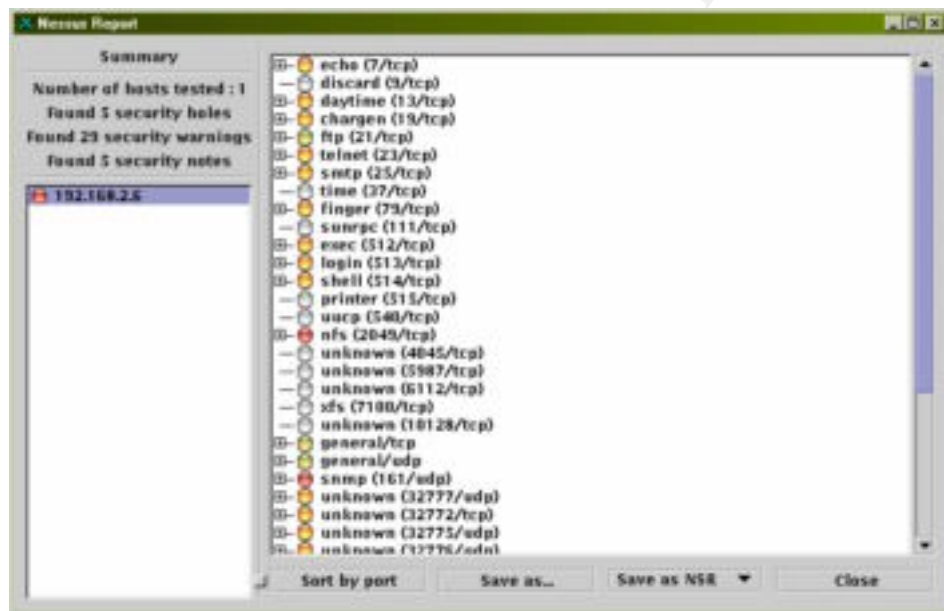


Figure 10 - Nessus scan results

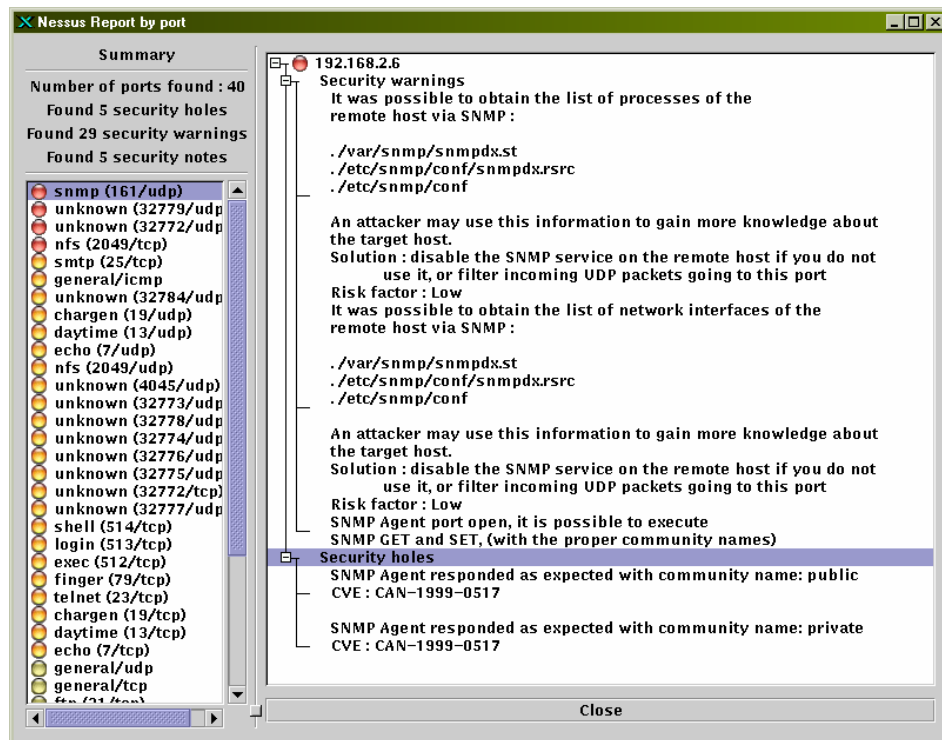


Figure 11 - Nessus results detail

It is expected that the initial scan s will take one day, however additional time will be required to manually confirm the results of the scan (e.g. if permitted, actually running exploit code against the exposed services).

1.4.1.4 Internal probing and verification of configuration

Internal probing of the network will start by using the same procedure as for the external scan, however this time the probes will be launched from inside the various zones of the network, including the Service Networks. In this case, logs will also be inspected to ensure that logging is occurring as expected.

Secondly, the configuration of the key components will be verified both by examining the configuration files, policies, etc. and by using sniffers (e.g. tcpdump – <http://www.tcpdump.org>) to verify the functionality e.g. to ensure that no traffic leaks through firewalls, all blocked traffic is logged, etc.

For example, for Cisco routers the difference of the current configuration from the default may be shown by issuing the following command in configuration mode:

```
# show config
```

Individual details of configuration may be found by querying those directly.

e.g. To check logging configuration:

```
# show logging
```

with expected output being similar to:

```
Syslog logging: enabled (0 messages dropped, 0 flushes,
0 overruns)
Trap logging: level debugging, 21 message lines logged
```

e.g. To check a specific access list configuration:

```
# show ip access lists 120
```

with expected output being similar to:

```
Extended IP access "list" 120
  access-list 120 permit ip 192.10.10.0 0.0.0.15 any (250 matches)
  access-list 120 permit ip 192.10.11.0 0.0.0.3 any (200 matches)
  access-list 120 deny ip any any log -input
```

This shows how many times each rule has been matched.

To check all access lists:

```
# show ip access lists
```

The tcpdump output below is a section of output taken during the Nessus scan shown above. In particular it shows part of the SNMP vulnerability scan and tcpdump's ability to perform a degree of application layer decoding to provide greater detail. It clearly shows us the scanning host (192.168.2.4) issuing SNMP GET (GetNextRequest) requests from the target (192.168.2.4) with a variety of common SNMP Community strings (C=Tivoli , C=community , etc.). For the Community strings shown below the failure of each request is also demonstrated (GetResponse(22) noSuchName):

```
17:17:36.728306 eth0 > 192.168.2.4.32774 > 192.168.2.6.snmp: C=tivoli
GetNextRequest(24) .1.3.6.1.2.1 (DF)
17:17:36.728306 eth0 < 192.168.2.6.snmp > 192.168.2.4.32774: C=tivoli GetResponse(22)
noSuchName@1 .1.3.6.1.2.1 (DF)
17:17:37.048306 eth0 > 192.168.2.4.32774 > 192.168.2.6.snmp: C=openview
GetNextRequest(24) .1.3.6.1.2.1 (DF)
```

```
17:17:37.048306 eth0 < 192.168.2.6.snmp > 192.168.2.4.32774: C=openview  
GetResponse(22) noSuchName@1 .1.3.6.1.2.1 (DF)  
17:17:37.418306 eth0 > 192.168.2.4.32774 > 192.168.2.6.snmp: C=community  
GetNextRequest(24) .1.3.6.1.2.1 (DF)  
17:17:37.418306 eth0 < 192.168.2.6.snmp > 192.168.2.4.32774: C=community  
GetResponse(22) noSuchName@1 .1.3.6.1.2.1 (DF)  
17:17:37.728306 eth0 > 192.168.2.4.32774 > 192.168.2.6.snmp: C=snmp  
GetNextRequest(24) .1.3.6.1.2.1 (DF)  
17:17:37.728306 eth0 < 192.168.2.6.snmp > 192.168.2.4.32774: C=snmp GetResponse(22)  
noSuchName@1 .1.3.6.1.2.1 (DF)  
17:17:38.038306 eth0 > 192.168.2.4.32774 > 192.168.2.6.snmp: C=snmpd  
GetNextRequest(24) .1.3.6.1.2.1 (DF)  
17:17:38.038306 eth0 < 192.168.2.6.snmp > 192.168.2.4.32774: C=snmpd GetResponse(22)  
noSuchName@1 .1.3.6.1.2.1 (DF)  
17:17:38.348306 eth0 > 192.168.2.4.32774 > 192.168.2.6.snmp: C=Secret C0de  
GetNextRequest(24) .1.3.6.1.2.1 (DF)  
17:17:38.348306 eth0 < 192.168.2.6.snmp > 192.168.2.4.32774: C=Secret C0de  
GetResponse(22) noSuchName@1 .1.3.6.1.2.1 (DF)
```

This phase will take two weeks if a thorough examination of the architecture proposed in this document is required.

1.4.1.5 Presentation of results and recommendations

The final phase of any audit is the presentation of the results to both the technology organisation and the business leads.

The presentation will cover:

- The existence of security policies, procedures, etc .
- Their relevance to the business objectives
- The mechanisms used to enforce them and their effectiveness
- The results of the external probe i.e. how big a risk is the network exposed to from outside?
- The results of the internal probe i.e. how easily could the network and business information be compromised from inside?
- Relevant corrective measures
- Ongoing maintenance of security and good practice

The presentation will focus on both the good and the bad aspects of the audit results. Particular effort will be made to ensure that there is no feeling that all is already lost.

The presentations, and detailed discussion of the results with the relevant parties are likely to take at least 2 days.


1.4.2 Financial Cost

Due to the use of publicly available software, the cost of the audit itself is mainly due to the man hours consumed. Assuming a UK daily consultancy rate of £1200 the audit described here could cost in excess of £20000. However, potential regulatory issues aside, there is no reason that much of this process could not be carried out in house. Engaging a third party does have the advantage though that they are often perceived by senior management to be more objective and the cost often forces serious consideration of the results.

1.4.3 Suggested Technical improvements

1.4.3.1 Resilience

The most obvious deficiency at the perimeter of the network is the lack of redundancy. This could be alleviated by having two ISP connections and redundant Nokia firewalls (via VRRP) e.g.



Both Nokia firewalls would connect to each service network and have single shared virtual router IP address.

1.4.3.2 Security

Service network services could be made more resilient by utilising clustering software (e.g. StoneBeat) or load balancers (e.g. Alteon ACEdirector 4).

These options, however, are expensive and there would need to be strong justification to the business before they could be implemented.

Service protection is strong, particularly due to the extensive use of proxies within the architecture. If greater protection is required, however, products such as Entrust's getAccess could also be used to enable granular access control to the web application.

1.4.3.3 Forensics


In order to strengthen our ability to react after an attack the following could be implemented:

- Full traffic capture of all traffic to and from the service networks. Results to be held for a defined period
- Audit logging on the Solaris hosts. The Solaris Basic Security Module can be configured to log all activity on the host.

Both of these would provide a huge amount of information in the event that analysis had to be performed, but both also require large amounts of storage. In addition, there may be legal ramifications to the monitoring and storage of such activity.

2. Assignment 4 - Design Under Fire

The aim of this section is to identify the weaknesses of a particular design solution. The design that will be considered here is that produced by Ken Colson:



Some of the primary features of interest are

- Border router – Cisco 2600
- External firewall – Linux IPChains
- Internal firewall – Linux IPChains
- VPN device (Cisco VPN 5000) spans both firewalls
- IDS spans internal firewall
- Publicly accessible hosts - Linux

2.1 Social Engineering

Depending on the ultimate aim of any system compromise attempt, the easiest method to use may well be just to ring up and ask. If the names of users, customers, partners, etc. can be gleaned from publicity material, mailing lists or other sources then it these could be used in a phone call to extract more useful information (for instance payment/account information, personal details, account passwords, etc.).

2.2 Border Router

The primary features to be noted from the configuration of the border router are:

- There is no logging of access list matches

Any scan attempts using tools such as NMAP or Nessus directed at the border router will go completely unlogged.

- No remote logging at all
Log events are unlikely to be monitored in 'real time' thus likely providing a window of opportunity out of hours in which our actions will go unnoticed
- There is no indication of whether or not SNMP is configured – if it is then there is no access list to control access to it so SNMP services may be freely probed using brute force methods to compromise the community string (e.g. using Nessus). Once the community string is determined then standard SNMP tools such as NET-SNMP (<http://net-snmp.sourceforge.net/>) may be used to read and, potentially, manipulate router configuration. Such information is likely to provide highly valuable during any further intrusion attempt
- HTTP server – the HTTP administration server is not disabled and it is not protected by an access list. Cisco IOS versions up to 12.2x are vulnerable to arbitrary administrative access via this service (Cisco IOS HTTP Configuration Arbitrary Administrative Access Vulnerability - <http://www.securityfocus.com/bid/2936>). The router may be compromised using a standard web browser.
- There is no egress filtering

2.3 Perimeter Firewall

- IPChains is not stateful

If we can compromise a protected machine the firewall will not block any outgoing traffic we generate as long as we set our source port appropriately.

In addition, the DNS rule (`ipchains -A input -i $EXT_INTERFACE -p udp -s $ANYWHERE $UNPRIVPORTS -d $OUR_NAMESERVER 53 -j ACCEPT`) is insufficient to allow proper DNS operation. It will have to be widened to allow replies from any port (many DNS servers send replies from port 53 and indeed the source port can be user configured). Replies may also be long replies necessitating allowing TCP inbound to the DNS server as well. When these changes are made to permit correct operation then the DNS server is much more exposed.

- All filtering is at the external firewall interface

If we can compromise a service network machine then we are free to attack the firewall from the reverse side with no firewall rules to limit what we can do. As the firewall is administered using OpenSSH we could attempt compromise via either of the following SSH vulnerabilities:

OpenSSH	PAM	Session	Evasion	Vulnerability	-
http://www.securityfocus.com/bid/2917.html					
OpenSSH		UseLogin		Vulnerability	-
http://www.securityfocus.com/bid/1334.html					

Further system vulnerabilities may also be discovered by probing the firewall (e.g. NMAP, Nessus)

2.4 Exposed Services

2.4.1 SMTP mail

Email service is being provided by Sendmail. No indication is given as to how Sendmail has been secured so the following should be tested:

- Mail relaying
- Flooding local accounts with mail e.g. sending email to the root user on that host
- Mail bomb attack on sendmail 8.8.x
<http://www.hack.co.za/download.php?sid=1310> (HELO hole) safebomb.c
- Mail flooding on sendmail 8.9
<http://www.hack.co.za/index.php?page=content&osid=237> against.c

These will not render local access, but will create severe inconvenience and, in the case of relaying, potential damage to reputation (and addition to anti-relaying block lists e.g. <http://mail-abuse.org/rbl/>)

2.4.2 Web Server

It is assumed that the web server in use will be Apache. Historically Apache has proven to be quite secure, though vulnerable to mis-configuration:

- Ability to view CGI scripts -
http://www.suse.com/de/support/security/adv7_draht_apache_txt.txt
- Apache Artificially Long Slash Path Directory Listing Vulnerability (<http://www.securityfocus.com/bid/2503.html>) exploits on site

Often the best mode of attack is to target the application rather than the server. For instance, many applications have vulnerable CGI scripts (e.g. WWWBoard Password Disclosure Vulnerability - <http://www.securityfocus.com/bid/649.html>) that can be discovered manually or using automatic tools such as Whisker (by Rain Forest Puppy) - <http://www.wiretrip.net/rfp/p/doc.asp?id=21&iface=2>

2.4.3 DNS Server

DNS has been the source of a wide range of critical vulnerabilities e.g.

- Multiple Vendor BIND (NXT Overflow & Denial of Service) Vulnerabilities (<http://www.securityfocus.com/bid/788.html>) exploit <http://www.hack.co.za/daemon/named/t666.c>
- ISC Bind 8 Transaction Signatures Buffer Overflow Vulnerability (<http://www.securityfocus.com/bid/2302.html>, <http://www.cert.org/advisories/CA-2001-02.html>) remote exploit - <http://www.hack.co.za/download.php?sid=1187>
- DNS Cache poisoning - http://www.sans.org/infosecFAQ/firewall/DNS_spoof.htm

By default ISC bind performs recursive queries. By forcing the server to query another server under our control we can supply new name/address mappings that will be cached for future use. In this way we can redirect outgoing traffic. This could be used to intercept outgoing email, or direct web connections to well known sites so we can serve up malicious web pages (e.g. Microsoft IE and OE XML Stylesheets Active Scripting Vulnerability - <http://www.securityfocus.com/bid/2633.html> or Microsoft Windows Media Player .WMZ Arbitrary Java Applet Vulnerability - <http://www.securityfocus.com/bid/2203.html>)

2.4.4 VPN Server

Although I have been unable to find any specific exploits for the VPN concentrator, it should be remembered that you don't have to compromise the concentrator in order to compromise the VPN. The target of any attack will be the systems which are connecting to the concentrator. Once one of these is compromised then the whole internal network beyond the concentrator will be exposed. Combined with the lack of firewall restricting traffic from VPN (and hence high degree of trust) this is a very serious vulnerability.

2.5 Intrusion Detection System

The NFR IDS spans the internal firewall and hence provides an alternative route into the network if it can be compromised. If it can be disrupted then this would also reduce the possibility that we would be detected:

- NFR (<http://www.securityfocus.com/bid/63.html>)

Upon receiving a IP packet with the protocol field set to TCP but with an all null TCP header and data section nfrd will die. nfrd will be automatically restarted but the attack packet does not get logged. The nfrd.log file will also be overwritten by the new instance of nfrd. This opens a window of opportunity for an attacker to send packets that will not be processed by nfrd while it is restarting.

- Buffer overflow in webd in Network Flight Recorder (NFR) 2.0.2 - Research allows remote attackers to execute commands. (Name CVE-1999-0375 - <http://cve.mitre.org/cgi-bin/cvekey.cgi?keyword=CVE-1999-0375> , <http://www.paraproduct.com/Advisories/Guests/1999/990209.txt>)

2.6 Internal Firewall

- Linux IPChains (non -stateful)
- Masquerading (Network Address Translation) for outgoing traffic

The base of the ruleset is:

```
ipchains -A input -i $INT_INTERFACE -d $ANYWHERE -j ACCEPT
ipchains -A output -i $INT_INTERFACE -d $ANYWHERE -j ACCEPT
ipchains -A forward -i $EXT_INTERFACE -s $INT_NETWORK -j MASQ
ipchains -A input -i $EXT_INTERFACE -p udp \
-s $ANYWHERE $UNPRIVPORTS -d $OUR_NAMESERVER 53 -j ACCEPT
ipmasqadm portfw -a -P upd -L $OUR_EXTERNAL_ADDRESS 514 -R \
$OUR_SYSLOG 514
ipchains -A input -i $EXT_INTERFACE -p udp -s $OUR_MACHINES 514 \
-d $OUR_SYSLOG 514 -j ACCEPT
```

This allows any and all traffic out of the environment i.e. the Service Network is highly exposed in internal activity (e.g. if our social engineering work can yield a rogue dial-in account, etc.)

Alternatively, if we can compromise the Service Network, then we could attack the weakly protected internal DNS (all udp allowed) using the exploits above or try to compromise the protection provided by the masquerading using:

Access Validation Error - <http://www.securityfocus.com/bid/1078> - which allows remote re-writing of UDP masquerading rules.

Appendix A. References

The following references have been used in the preparation of this document:

- [1] *Building Internet Firewalls*. Chapman & Zwicky.
- [2] *Managing Cisco Network Security*. Wenstrom. 2001
- [3] *TCP/IP Illustrated, Volume 1*. Stevens
- [4] *Network Intrusion Detection. An Analysts Handbook*. Nortcutt. 1999
- [5] *TCP/IP for Firewalls and Intrusion Detection*. SANS Institute
- [6] *Firewalls 101: Perimeter Protection with Firewalls*. SANS Institute
- [7] *Firewalls 102: Advanced Perimeter Protection and Defense*. SANS Institute
- [8] *VPNs and Remote Access*. SANS Institute
- [9] *RFC 1918 - "Address Allocation for Private Internets"*. Rekhter, Moskowitz et al. February 1996 (<http://www.faqs.org/rfcs/rfc1918.html>)
- [10] *RFC 1700 - "Assigned Numbers"*. Reynolds & Postel. October 1994 (<http://www.faqs.org/rfcs/rfc1700.html>)
- [11] *RFC 2018 - TCP Selective Acknowledgment Options*. Mathis, Mahdavi et al (<http://www.faqs.org/rfcs/rfc2018.html>)
- [12] *RFC 2402 - IP Authentication Header*. Kent & Atkinson (<http://www.faqs.org/rfcs/rfc2402.html>)
- [13] *DNS and BIND*. Albitz & Liu. 3rd Edition. 1998
- [14] *RFC 2136 - Dynamic Updates in the Domain Name System (DNS UPDATE)*. Vixie, Thomson et al (<http://www.faqs.org/rfcs/rfc2136.html>)
- [15] *VPN 3000 Concentrator Series User Guide*. Release 3.0. Cisco. March 2001
- [16] *Juniper Firewall*. (<http://www.obtuse.com/juniper/>)
- [17] *Security Focus* - <http://www.securityfocus.com>
- [18] *Wiretrip* - <http://www.wiretrip.net>
- [19] <http://www.hack.co.za/>
- [20] *The SANS institute* – <http://www.sans.org/>

End of Document