



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.

Table of Contents	1
Scot_Hartman_GCFW.doc	2

© SANS Institute 2000 - 2002, Author retains full rights.

SANS GCFW Practical Assignment v1.6

**SANS Rocky Mountain
Denver, CO**

July 2001

Prepared by:

Scot Hartman

August/September, 2001



Table of Contents

1.0	ASSIGNMENT 1 - SECURITY ARCHITECTURE ..	4
1.1	SOLUTION ASSUMPTIONS.	4
1.1.1	<i>Basis for Security Architecture Deployment</i>	4
1.1.2	<i>Existing Architecture</i>	4
1.1.3	<i>Redesign Strategy</i>	5
1.2	WRITTEN SECURITY POLICY	7
1.3	HARDWARE AND SOFTWARE USED IN THE DESIGN	8
1.3.1	<i>Firewalls</i>	8
1.3.2	<i>VPN Appliances</i>	11
1.3.3	<i>Routers</i>	12
1.3.4	<i>Intrusion Detection</i>	12
1.3.5	<i>Miscellaneous</i>	12
1.4	PROPOSED SECURITY ARCHITECTURE	13
1.5	SECURITY ARCHITECTURE EXPLAINED	14
1.5.1	<i>The Border</i>	14
1.5.2	<i>The Customer Networks</i>	16
1.5.3	<i>The Partner/Reseller Networks</i>	18
1.5.4	<i>Web & Database Management Network</i>	19
1.5.5	<i>Security & Network Management Network</i>	20
1.5.6	<i>Management VPN Network</i>	21
1.5.7	<i>Mirror Site / Disaster Recovery</i>	21
1.5.8	<i>Corporate HQ / Offices</i>	22
1.5.9	<i>High-availability Components</i>	22
1.6	DEPLOYMENT PRIORITIES	23
2.0	ASSIGNMENT 2 - SECURITY POLICY.	24
2.1	ASSUMPTIONS	24
2.2	BORDER ROUTER	24
2.2.1	<i>Overview</i>	24
2.2.2	<i>Detailed steps to set ACLs & harden system</i>	24
2.3	EXTERNAL (PRIMARY) FIREWALL	28
2.3.1	<i>Overview</i>	28
2.3.2	<i>Detailed steps used to configure Nokia appliance</i>	28
2.3.3	<i>Detailed steps used to configure Checkpoint and to set the security policy</i>	31
2.3.4	<i>External (Primary) Firewall Security Policy Summary</i>	42
2.4	LAN-TO-LAN VPN APPLIANCES	43
2.4.1	<i>Overview</i>	43
2.4.2	<i>Steps used to configure and deploy the Nokia VPN appliances</i>	43

3.0	ASSIGNMENT 3 - AUDIT OF SECURITY ARCHITECTURE .	47
3.1	PLANNING THE SECURITY AUDITS	47
3.1.1	<i>Pre-deployment Audits</i>	47
3.1.2	<i>Deployment Audits</i>	48
3.1.3	<i>Ongoing Audits</i>	50
3.2	CONDUCTING THE SECURITY AUDITS	50
3.2.1	<i>Pre-deployment Audits</i>	50
3.2.2	<i>Deployment Audits</i>	58
3.2.3	<i>Ongoing Audits</i>	61
3.3	ANALYZING THE SECURITY AUDITS	61
3.3.1	<i>Pre-deployment Audits</i>	61
3.3.2	<i>Deployment Audits</i>	65
3.4	CORRECTING VULNERABILITIES DISCOVERED DURING THE SECURITY AUDITS	65
4.0	ASSIGNMENT 4 - DESIGN UNDER FIRE .	66
4.1	ARCHITECTURE UNDER FIRE	66
4.2	ATTACK AGAINST THE FIREWALL ITSELF.	67
4.3	DENIAL OF SERVICE ATTACK.	69
4.4	ATTACK PLAN TO COMPROMISE AN INTERNAL SYSTEM	70
4.5	CONCLUSION	70
5.0	ACKNOWLEDGMENTS AND REFERENCES	71
	APPENDIX A	73

1.0 ASSIGNMENT 1 – SECURITY ARCHITECTURE

1.1 SOLUTION ASSUMPTIONS

GIAC Enterprises (GIACE) is described as a growing Internet startup and therefore relies heavily upon its ability to provide an effective and secure web presence. The expanding company, having just completed an acquisition, is poised to dominate the fortune cookie market. This position is expected to earn the company \$200 million per year in on-line sales. With this potential market position and the monetary rewards possible, the company cannot afford to jeopardize its future on ill-conceived or half-baked security architectures.

1.1.1 Basis for Security Architecture Deployment

With the rapid growth of the company and the recent acquisition, GIACE is seeking outside assistance to help them develop a secure architecture. In addition to acquiring the other company, GIACE has also recently received an influx of venture capital on the strength of its market projections and its past performance. This puts them in the position to develop an architecture that is both secure and flexible for their continued growth.


Additional incentive comes from recent, highly publicized attacks on a similarly sized Internet company. This company had their web site repeatedly defaced and their customer database violated. They were unable to recover fully from the loss of consumer and investor confidence and are rumored to be considering bankruptcy protection.

1.1.2 Existing Architecture

As an existing company, GIACE already has an IT staff and an Internet presence. The existing network and security architecture, however, were designed during the company's infancy when the budget was very tight and the focus was on proof-of-concept for the product itself. As the company became more successful, some rudimentary security measures were undertaken, but security is still not an integrated part of the architecture.

The initial design consisted of a couple of servers plugged directly onto a publicly IP'd subnet. This subnet was delivered via a T-1 link from their ISP and the only protection was a hastily developed Access Control List (ACL) on the router. The ISP hosts GIACE's Domain Name Service (DNS). Eventually, as GIACE added more functionality and received more business, they added more servers. Once they were large enough, they relocated to an Internet Data Center (IDC) hosting facility. A firewall was eventually installed between their Internet router and the server network. The original T-1 link continues to be used by the employees to access the Internet and to remotely access their servers which are now located at the collocation facility. The firewall allows all connections from the corporate network.

The state of the architecture can be easily summarized as below...



GIACE is not an unusual case of a security-as-you-go scenario. They have implemented Band-Aid solutions as they've grown but never fully integrated security into their deployment. They currently have only a single layer of perimeter protection and have all their eggs stacked together into the single basket behind it. The corporate location currently is even less protected. The low level of protection for the corporate network and the trust imparted upon it by the e-commerce site's firewall allows corporate to become a serious Achilles heel to the security of their service.

Some realization of these vulnerabilities and a growing understanding of the need for a more robust solution have served as a catalyst for redesign. The integration of new partnerships, the growth of the industry, and their dominant market share within it have sparked an impetus to seek outside expertise.

1.1.3 Redesign Strategy

Review of the existing architecture shows several areas that will need improvement. Even with a capital infusion, however, it is important to prioritize the areas of focus. The proposed security architecture is designed to answer most of the vulnerabilities in the existing architecture while allowing for staged deployment and flexibility of implementation. The proposed design will also point to areas of further development as the company continues to grow.

The major areas of focus are...

- Separation of the functionality and vulnerability categories
- Layering of the Defense
- Secure Access for Customers / Partners

- Integration of Different Technologies
- Secure Access for Management / Administration
- Redundancy
- Intrusion Detection
- Additional Filtering / Protection for Denial of Service

Separation of the functionality and vulnerability categories

This is primarily separating servers or networks by what job they are designed to do. Putting Web servers on their own network. Putting Mail or DNS on their own network. Traditionally, each functional area tends to have its own vulnerabilities and separating them mitigates the potential damage if one is breached. In addition to separating by function, it is also a good idea to separate by levels of access: public servers vs. non-public servers on different subnets with a security enforcement point in between. This leads into Layering of the Defense.

Layering of the Defense

Use of multiple enforcement points instead of a single perimeter security scenario. The border router will conduct initial screening and a firewall to protect the public servers. Access of databases from the web servers will be through a back-end network with another firewall for protection. An additional level of security can be inserted between the development servers and master database servers. Use of a separate security network for out of band (OOB) management of security components for added layering and security. Part of the layering is the use of host-based tools to protect the servers themselves. Integrating TCP wrappers, file integrity checkers, etc. as well as ensuring code levels and patches are applied properly.

Secure Access for Customers / Partners

Allows for protected access that will both shield privacy and ensure only legitimate usage is performed.

Integration of Different Technologies

Using different technologies can add to complexity. However, use of various technologies and different vendor products at different points along the layered defense can allow one point to protect against vulnerabilities in another. The combination of a layered architecture and multiple technologies/techniques can increase the work required to break into the inner layers of protection. No protection is absolute, the objective is simply to raise the effort required to be higher than the potential reward.

Secure Access for Management / Administration

Corporate users who maintain or need access to the e-commerce servers will need a way to do this that is secure and minimizes the ability for some unauthorized elements to take advantage. Some may believe that the most secure way is to not allow any form of remote access, but this is not the magic bullet that it seems. Administration and access need to be both secure and workable. If security measures make it too hard for people to do their job, they will not take them seriously and will resist them. In fact, if it is too difficult, administrators will be less likely to perform proper patch updates and code upgrades. This may actually lower the overall security stance. The goal will be to make administration a livable compromise between convenience and security. The design will include secured remote access.

Redundancy

Refers to the addition of multiple connections to the Internet, deployment of high-availability firewalls/routers/switches, construction of a mirror location or disaster recovery site.

Intrusion Detection

Intrusion detection can be included to alert administrators of suspicious activity via host-based or network-based monitoring.

Additional Filtering / Protection for Denial of Service

There are many malicious attacks that can be launched against someone on the Internet. Some attacks can best be stopped at various points in the network architecture. As the security architecture matures, different tools can be integrated to begin to harden the site against the various forms of Denial of Service (DoS) attacks.

1.2 WRITTEN SECURITY POLICY

GIACE will need to develop a written security policy. This policy will fit into the hierarchy under a Corporate Policy and possible Local Policies that will govern specific company areas. Under these policies should be developed procedures for password usage, resource usage, incident handling, backups, etc. These need to be developed by the company and be written down so that they can be referenced.

As part of the security policy, there should be a written local security policy to cover the e-commerce site and provide guidance for the security architecture. This will be a written statement of what the company is trying to accomplish, what it will or will not use to accomplish these goals, and who is responsible for each area.

The preceding section outlines the redesign strategy. The following sections outline the proposed security architecture. The hardware proposed is detailed along with reasons for each choice. The architecture is laid out and then each portion/component is explained in the following details. These

outlines and reasoning will provide input and can serve as a draft for the final written local security policy.

Once the local security policy is completed. Each security enforcement point will have an individual specific policy that will outline its specific part of the overall security policy. These devices will then use these policies to shape their implementation and maintenance.

© SANS Institute 2000 - 2002, Author retains full rights.

1.3 HARDWARE AND SOFTWARE USED IN THE DESIGN

This section provides an explanation of some of the equipment that will be used in the proposed architecture. The focus is upon the network and security architecture and does not cover the requirements for application servers, database servers, etc.

1.3.1 Firewalls

Five major firewall types are proposed. Each type has individual strengths and weaknesses. All the proposed firewall types are commercially developed and supported. There are many fine firewall choices that are available through the open-source community but the decision to use commercial products is based upon the following criteria.

- The commercial products have an infrastructure of standardized training and certifications that will allow these objective criteria to assist in selection of support staff.
- The commercial products themselves have technical support staff and can be contractually obligated to provide critical assistance when needed.
- The commercial products have incentive to ensure that code revisions, patches, and updates are regularly available. With many competing vendors, and market share at stake, the need is apparent to stay ahead of vulnerabilities. The open-source community can typically react faster to provide patches, but there is risk of unevenness of application, trust in the source, the continued support of the application, or the skills needed to apply the provided support.
- There are typically major studies and comparisons between firewall products. With the monetary investment required for firewalls, there is traditionally an interest in commercial product comparisons and vulnerability studies. This equates to technical journals, online and print, using their resources to conduct the studies to attract readers. Commercial firewall products have typically fared well in demonstrated security and reliability. Open source tools will typically dominate the expense category for obvious reasons, but are not more secure on average.

The commercial products have comparable reliability and security. Their major advantages are in the training, support, known quantity, and an available pool of qualified administrators. All firewalls will be managed out-of-band on a dedicated management network that will be non-routable to any network and inaccessible from any other network connected to the firewalls.

The following firewalls will be used.

- Cisco Router ACL

The lowest level of firewalling, access control lists perform packet filtering and will allow the border router to do an initial screening of the traffic passing through them. This provides an added layer to defense in depth.

This design will use the existing Cisco 7206 router that the company already owns. The configuration will be revised, access control lists will be added, and the IOS level will be updated to v12.2.

- Checkpoint Firewall-1 Firewall

Checkpoint is a software company that creates firewall software to run on various vendor operating systems. Firewall-1 is stateful-inspection firewall that adds to packet filtering by including the ability to look deeper into packets and track the state of a permitted connection. Nokia is a hardware vendor that creates network appliances designed to host Checkpoint firewall software. Nokia uses a hardened version of BSD Unix called IPSO and, like Checkpoint, provides regular patches and updates. The existing firewall was a Checkpoint firewall running on a Windows[™] NT system. The license will be reclaimed and used.

This design will utilize Firewall-1 v4.1 SP3 loaded on the Nokia IP330 hardware platform using IPSO 3.3. As the company grows, Nokia offers higher models in the IP line for added throughput and more DMZs.

- Cisco PIX Firewall

Cisco is a network hardware company that also creates the software to operate their appliances. The security product suite from Cisco includes the PIX firewall line that provide stateful-inspection on a lean, hardened, hardware-accelerated appliance. The Cisco firewalls typically have an advantage over the Checkpoint firewalls in performance vs. price but the management of them is not as polished or user-friendly. The Cisco PIX is an excellent firewall where high performance is desired and a fairly simple rule base is needed.

This design will utilize the Cisco PIX 515-UR firewall utilizing v6.01 PIX-OS. As the company grows, Cisco offers higher models in the PIX line.

- **Netscreen Firewall**

Netscreen is another network hardware company that also creates the software for their appliances. The Netscreen boasts high performance due to use of ASICs in the design and is another stateful-inspection firewall. The Netscreen will be used because of a unique ability to act as a true firewall at layer 2 (OSI model). While most firewalls act as routers (or higher), the Netscreen can behave as a true switch in what is called 'transparent mode'. While some other vendors can simulate a switch by using proxy ARP, this still leaves the firewall's MAC address exposed. The Netscreen, however, will allow a subnet to be divided invisibly and will simply not forward denied packets.

The design will utilize the NS-10 running ScreenOS v2.6 where some added firewalling might be needed. As the company grows, Netscreen offers higher models in its line.

- **Raptor Firewall**

Raptor is a software company, like Checkpoint, that creates firewall software to run on a third-party operating system. Unlike Checkpoint, however, Raptor is an application layer firewall that terminates connections destined through it and acts as a proxy for the protected device. This type of firewall must have the capability to, not only inspect the traffic, but to act as a client and a server for the desired type of traffic. While the typical web server daemon has traditionally been designed for functionality without looking to how a malicious user could misuse it, proxy programs in application layer firewalls are. The firewall's web server daemon can be designed to handle typical malicious attacks that would cripple the actual server, and pass only legitimate traffic through. A packet filtering or stateful inspection firewall does not offer this degree of defense. This added protection does not come without a price however. Proxy firewalls are typically slower than their counterparts and are limited in the services they can support by the proxies available from the vendor.

This design will utilize Raptor firewall v6.5. Although Raptor originally utilized Windows operating systems, they now support Solaris. Solaris 2.8 will serve as the operating system for the firewall.

1.3.2 VPN Appliances

Two types of virtual private networks (VPNs) will be utilized in the proposed design: remote user VPN for partners and resellers and a point-to-point VPN mesh for administration of the e-commerce site and management of the network. The VPNs will be built utilizing dedicated VPN appliances.

There are two major reasons to use division of labor between the firewall and VPN devices: encryption is a highly processor intensive task and modularizing makes upgrading easier.

The firewall's main task should always be to inspect the traffic that is attempting to pass through it. The high level of processing power required for encryption and the regular key exchanges for VPN tunnels make it easy to self-DoS. If your e-commerce firewall doubles as your VPN terminator, the VPN tunnel has the potential to utilize all the available resources and prevent the traffic that may be the bread-n-butter of the organization.

A firewall that may be powerful enough to do its primary job may be underpowered to handle the encryption duties of the VPN. Upgrading the firewall software and hardware is typically an expensive task and, once upgraded, may need to be performed again as VPN traffic and e-commerce traffic increases. Additionally, the add-on encryption cards for firewalls often cost as much, or more, than the stand-alone appliance.

VPN appliances are relatively inexpensive, reliable, and can be sized independent of the firewall. The VPN appliance can offload the encryption duties from a firewall and can grow as the need grows without affecting the firewall (very important if you don't like having outage maintenance windows). VPN appliances are also specialized to perform one task, and perform it very well. They are typically much easier to configure and maintain than to force this function onto the firewall and also allow the firewall's configuration to be simpler as well. Complexity of configuration is typically anathema to security; overly complex makes it more difficult to troubleshoot and to dissect.

The devices used in the design.

- Remote user VPNs
This will be provided using the Cisco 3000 VPN Concentrator that will terminate VPN tunnels from client software loaded on partner/reseller machines and accessed through the Internet. The device will enforce IPsec.
- Point-to-Point VPNs
This will be provided using the Nokia Crypto Cluster 500 VPN appliance. The Nokia allows up to 500 devices to be fully meshed using IPsec tunnels and will be used over the Internet (then later Frame Relay) to link various e-commerce locations and corporate HQ on a protected management network.

1.3.3 Routers

The routers used will be the ones previously owned by GIACE. The routers are Cisco 7106 series and will use version 12.2 IOS.

1.3.4 Intrusion Detection

Intrusion detection systems (IDS) can be deployed as part of the security architecture. There are two major flavors of IDS: host-based and network-based. IDS are designed to notify, or act upon, suspicious behavior and can be a valuable addition to the security of the design. IDS is not, however, a replacement for a good perimeter defense. They are used to augment the security architecture by providing information about traffic that may have already breached the defenses or is suspicious in nature. IDS are typically expensive to install and take special training to support.

The initial focus of this proposal is to build the layering of the perimeter defense. The design will outline where network IDS may be deployed but is not part of the initial rollout. Host-based integrity tools such as Tripwire should be installed on servers and network IDS should be looked into as the security architecture matures.

Commercial vendors such as Cisco or ISS provide IDS solutions, some companies such as Counterpane or ISS provide managed solutions. The open-source community has many good IDS products that can be utilized such as Snort by Marty Roesch.

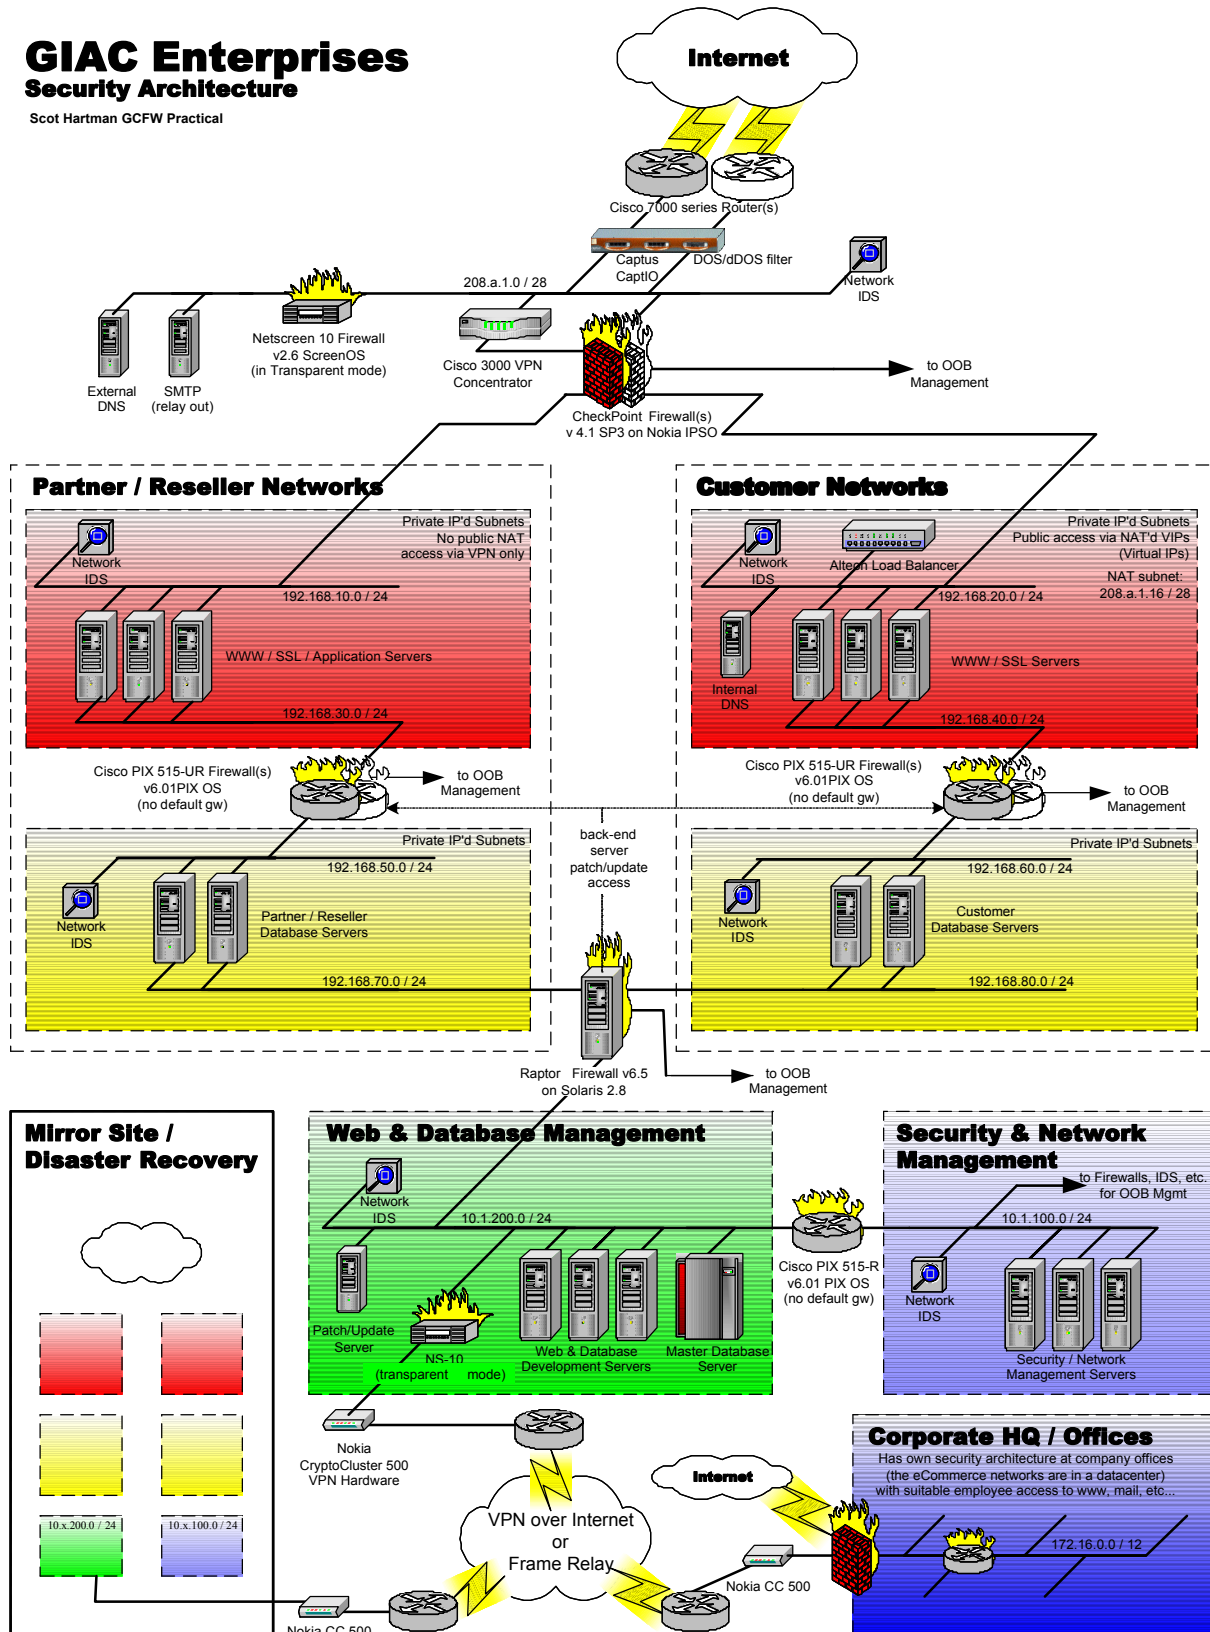
1.3.5 Miscellaneous

Captus is a company that has incorporated filtering into a high-throughput device that can be used to help protect against various sources of flooding DoS attacks. This is functionality that can be put upon other devices in the network, but the Captus device does an exceptional job of using thresholds and is very flexible. This is an added device that can take some of the load off of the border router and the firewalls and falls into the philosophy of defense in depth. This device can be deployed after the rest of the network is in place if needs determine its value. It is included in the proposed design to show where it can best be implemented to heighten its value.

1.4 PROPOSED SECURITY ARCHITECTURE

GIAC Enterprises Security Architecture

Scot Hartman GCFW Practical



1.5 SECURITY ARCHITECTURE EXPLAINED

The proposed security architecture is broken down into several discrete sections. Each section will be discussed separately along with the components that are contained in them. Additional components that will augment the entire architecture as a whole will be discussed after the individual sections.

The general sections represented in the security architecture are:

- The Border.
- The Customer Networks. Consisting of four networks in two basic halves: public facing and back-end processing.
- The Partner/Reseller Networks. Consisting of four networks in two basic halves: public facing and back-end processing.
- Web & Database Management Network.
- Security & Network Management Network.
- Management VPN Network.
- Mirror Site / Disaster Recovery
- Corporate HQ / Offices.
- High-availability Components

1.5.1 The Border

Between the Internet and the first level of firewalls is the Border. This may seem like a fairly unimportant piece from a security perspective, but it can be used to meet some critical needs. This subnet will also probably be an early one to be fielded with a network based IDS when the architecture matures to that point.

Border Router(s)

The border router will be performing initial traffic filtering using access control lists. The primary focus will be to screen out IP addresses that should never be legitimately seen traversing the Internet. The list of recommended IPs to screen can be found at SANS (<http://www.sans.org/dosstep/index.htm>). The summary of the list is found below.

0.0.0.0/8	Historical Broadcast
10.0.0.0/8	RFC 1918 Private Network
127.0.0.0/8	Loopback
169.254.0.0/16	Link Local Networks
172.16.0.0/12	RFC 1918 Private Network
192.0.2.0/24	TEST-NET
192.168.0.0/16	RFC 1918 Private Network
224.0.0.0/4	Class D Multicast
240.0.0.0/5	Class E Reserved
248.0.0.0/5	Unallocated
255.255.255.255/32	Broadcast

The border router will also be used to screen for spoofing attacks. The router will not allow any traffic inbound on its Internet interface that has a source using an IP of an internal system and will also prevent anything out to the Internet that does not have a source of a valid internal IP.

The border router will also follow SANS guidance to prevent IP directed broadcasts to prevent the site's use as a broadcast amplification point. The command "no ip directed-broadcast" is default on Cisco IOS 12.0 or higher but will still be an auditable item.

Spare Subnet

An additional IP subnet will be provisioned and routed to the border. This subnet will not be advertised in DNS, will be non-contiguous from the subnet in use, and will remain unused until needed. This subnet's function is to serve as an alternate, publicly unknown, publicly routable, address space that can be cut over to if needed.

There are several uses for this subnet. It could be used for development, demonstration, or testing purposes. From a security perspective, it can be used as an alternate address space that can be shifted to if the primary subnet is subjected to any long-term DoS attacks. This was a technique used by www.whitehouse.gov to avoid the intended affects of Code Red.

Captus Filter

The Captus device has the ability to set threshold filtering that can assist in protecting against some forms of flooding DoS attacks. Although similar filtering functions can be implemented in either the routers or the firewalls, the Captus can do a better job of this and can be added under the principle of defense in depth. The filtering uses adjustable thresholds that can dynamically filter excessive types of traffic by individual source and destination IP. The threshold filters will be set only after normal traffic patterns are analyzed.

1.5.2 The Customer Networks

The customer portion of the architecture is made up of four subnets. These subnets are broken down into a public access subnet, a web server back-end net, a customer database subnet, and a database server back-end net. The subnets are divided by three firewalls: Checkpoint Firewall-1 firewall protecting the public access subnet, Cisco PIX firewall between the web servers and the database server, and a Raptor Firewall between the customer database servers and the web & database management subnet.

Customer Public Access Subnet

The customer public access subnet is a private IP addressed subnet. Access to the servers is via NAT'd virtual IPs on the load balancer. This can accomplish a few things from a security perspective. A specific individual server cannot be accessed directly from the Internet. Access is to a random server, based upon the load balancing algorithm. Bringing down any individual server will cause it to be taken out of the load-balance rotation, which can protect the overall accessibility of the customer site. Finally, Alteons are ASIC-based load balancing hardware that traditionally do a much better job of clearing their connection tables than the servers behind them. This can help protect against certain DoS attacks.

The company sends order confirmation to customers via e-mail when requested. In order to accomplish this, split DNS is utilized. An internal DNS server is located on the public access subnet and draws from the external DNS server located behind a Netscreen firewall on the border subnet. The internal DNS server handles DNS queries from the web/application servers. Mail is sent out via the relay SMTP server protected by the Netscreen firewall on the border subnet.

The Checkpoint Firewall on the Nokia platform will serve as the default gateway for the customer access servers. The firewall will employ both ingress and egress filtering. Inbound allowed connections will be limited to the ports needed for web and secure web traffic (TCP ports 80 and 443) and permitted to the virtual IPs only. DNS will be allowed between the internal and external DNS servers. The internal DNS server will have only a private IP assigned and will not have any public addressable access (no NAT address). SMTP will be allowed outbound only and will also be limited by source and destination between the DNS servers. The firewall itself will only accept management traffic to itself from the out-of-band management network. No other access will be permitted. For more information on management traffic to the firewall, see the section on Security and Network Management (Section 1.5.5).

Consistent with defense in depth standards, the servers themselves will be recommended to be equipped with TCP wrappers, file integrity programs, and kept up to date on code versions and patches. The same precautions will be taken on all servers to follow.

This subnet will probably be the first one to be fielded with a network based IDS when the architecture matures to that point.

Customer Web Server Back-end Subnet

The web/application servers will need to access the customer database servers to process order requests. This will be done via a back-end private IP network that connects through a Cisco PIX firewall. The servers will not have IP forwarding enabled. Use of the back-end network will serve several functions. First, access of the databases will require the use of additional ports and protocols. The use of the back-end network will prevent this sensitive traffic from traversing the public access network and will greatly limit the possibility of a clever attacker gaining access to it. Second, along the same vein, this will allow the TCP wrappers to be significantly tightened on the servers' external interfaces. Third, accessibility of the database servers will be more difficult for an attacker since any potential routing chain is broken. Finally, access from the web and database management subnet (for development upgrades, updates, patches, etc.) will also not need to traverse the public access network.

Customer Database Subnet

The customer databases will reside on a private IP'd subnet behind a Cisco PIX firewall. The firewall will be configured to only allow database specific traffic between the database subnet and the web server back-end subnet. The firewall will be kept 'dumb' for routing purposes. There will be no default gateway assigned and it will only know about the directly connected subnets and specific static routes. The firewall itself will only accept management traffic to itself from the out-of-band management network. No other access will be permitted. For more information on management traffic to the firewall, see the section on the Security and Network Management (Section 1.5.5).

Customer Database Back-end Subnet

This subnet will serve much the same purpose as the web/application server back-end subnet. (Again, the servers will not have IP forwarding enabled.) It will be a private IP'd subnet that will allow the master database to update and backup the customer database. Traffic needed for development upgrades, updates, patches, etc. will also use this back-end connection to prevent this traffic from traversing the customer-facing network. Protection between the back-end network and the web and database management network is via a Raptor firewall that is described more in the Web and Database Management section (Section 1.5.4).

1.5.3 The Partner/Reseller Networks

Like the customer portion of the architecture, the partner/reseller portion is made up of five subnets. These subnets are broken down into a VPN access subnet, a web server back-end net, a partner/reseller database subnet, a database server back-end net, and a remote user VPN subnet. The subnets are divided by three firewalls: Checkpoint Firewall-1 firewall protecting the VPN access subnet, Cisco PIX firewall between the web servers and the database server, and a Raptor Firewall between the customer database servers and the web & database management subnet.

Partner/Reseller VPN Access Subnet & Remote User VPN Subnet

The partner/reseller access subnet is a private IP addressed subnet. Access to this subnet is only available through a remote user VPN. The partner/resellers will connect using client software to a Cisco VPN concentrator that has a public IP address on the border subnet. That VPN appliance is responsible for user authentication, maintaining the VPN tunnels, and encryption/decryption. The VPN appliance has a direct crossover connection to the Checkpoint firewall on the Nokia platform. This will allow the firewall to limit the permitted traffic to the partner/reseller network without being hampered by the other functions carried out by the VPN appliance. This setup will also prevent traffic utilizing the VPN appliance from being able to bypass the firewall.

The purpose for the VPN access for the partner/reseller network is because they will require more powerful access into the system. They will be able to gain access via standard web-browser based tools (TCP ports 80 and 443) much like customers but with more available options. The partners/resellers will also use proprietary applications that run over TCP ports 9000 and 9001 to provide inputs to the system, check status, make queries, etc. It is protection of this communication that has prompted the use of VPNs for communication.

Like the customer public access network, the Checkpoint Firewall on the Nokia platform will serve as the default gateway for the partner/reseller VPN access network on a separate firewall DMZ. The firewall will employ both ingress and egress filtering. No traffic will be allowed between the customer and partner/reseller networks. This drop rule will reside at the top of the rule base to ensure it is matched first.

Inbound allowed connections will be limited to the ports needed for web and secure web traffic (TCP ports 80 and 443) and to the application TCP ports 9000 and 9001. As previously mentioned, the firewall itself will only accept management traffic to itself from the out-of-band management network. No other access will be permitted. For more information on management traffic to the firewall, see the section on the Security and Network Management (1.5.5).

As with the customer network subnets and consistent with defense in depth, the same server policies concerning patches, updates, file integrity checkers, etc. will be followed in

all the partner/reseller subnets.

This subnet will probably be one of the early ones to be fielded with a network based IDS when the architecture matures to that point.

Partner/Reseller Web/Application Server Back-end Subnet

This will be set up identically to the customer back-end subnet between the web servers and the database subnet. Different private IP addresses will be used, but the precautions and reasons are the same.

Partner/Reseller Database Subnet

This subnet and its components will also mirror the precautions and the reasoning already laid out in the customer database subnet. The partner and reseller databases are also separated by function. These databases, however, can provide inputs into the system via submissions from partners who created and input the fortune cookie sayings and resellers for sales and research updates. The purpose of the partner/reseller databases are to provide information to the partner/resellers, collect the inputs from the application servers, collate them, record them by source, and to transfer these to the master database server when it requests them.

Partner/Reseller Database Back-end Subnet

This subnet will also serve much the same function as the customer database back-end subnet. All the database traffic between the partner/reseller databases and the master database will take place over this subnet and will not traverse the same subnet as the partner/reseller-facing interfaces. All traffic between database servers is initiated by the master database at scheduled intervals. The methodology, protections, and reasoning are the same as the customer networks.

1.5.4 Web & Database Management Network

The web and database management network is a private IP addressed subnet that will house the web and database development servers and the master database servers. The subnet is protected from the various back-end subnets by the Raptor firewall. This firewall will have no default gateway defined and will only be aware of directly connected subnets and needed static routes. The firewall will provide ingress and egress filtering and will use available proxies. No traffic will be allowed between any of the back-end subnets. No traffic will be allowed to be initiated inbound to the web and database management network. The only allowed traffic will be specified traffic initiated from the web and database management network described below.

The master database server's function is to periodically draw inputs from the partner database,

provide inventory control, and to update the customer and reseller databases with available products. It is also responsible to collect billing information, provide for a central repository, and serve as a backup medium. The master database will be regularly backed up, along with other critical components, and stored off-site. (For more on this, refer to the Mirror Site / Disaster Recovery section 1.5.7.) The master database server will periodically initiate connections through the firewall to the slave databases to perform its described functions.

The web and database management subnet will provide private, back-end access for updating, patching, and managing the customer/partner/reseller server infrastructure. This traffic will be controlled by the Raptor firewall, which will allow traffic from the patch/update server to the back-end subnets.

The web and database management subnet will also provide private-side monitoring of processes on the servers. Monitoring servers will be allowed to poll the servers through the firewall.

Access to the web and database management subnet will be provided over a LAN-to-LAN VPN mesh using the Nokia CryptoCluster VPN appliance. (See section 1.5.6 on the Management VPN Network.) Behind the VPN appliance will be a Netscreen firewall. This follows a principle that all traffic must pass through a firewalled enforcement point, even if it is from a 'protected' VPN network. Defense in depth is one reason. The other reason is that this network is the crown jewel and is deserving of extra measures of protection.

1.5.5 Security & Network Management Network

The security and network management network provides out-of-band access and management for the firewalls, IDS sensors, etc. It is a separate subnet from the web and database management network and is protected by a Cisco PIX firewall. This network houses the necessary management and logging servers needed to support the security architecture.

The reason for an out-of-band management network is to limit, as much as possible, any potential vulnerability of the security appliances themselves. To an attacker, the firewall is itself a prime target for attack. If the firewall itself can be compromised, all the networks that it was installed to protect are now basically left on their own. Although it is important to keep abreast of the vulnerabilities of your firewall and keep them updated. It is just as important to minimize the potential vulnerabilities through secure architectural design if possible.

Exploits that allow attackers to gain access through firewalls are found regularly. Of these, most can be minimized by properly implementing ingress and egress filtering, keeping servers patched (especially on exploits that affect the ports you use), and through layered architecture. These exploits can range from mild to serious.

Exploits that can gain access or DoS the firewall itself, however, all potentially have catastrophic results. These are also periodically discovered. One thing to note is that vulnerabilities may be

available to the attacker community for a while before they become public knowledge (the White Hats are not omniscient). Patches may take a while to develop. A common vein for these vulnerabilities is the fact that most of them involve exploitation of management connections and any traffic allowed to the firewall itself. The best proactive defense to help eliminate this Achilles heal is to architect out any need for the public or protected subnets to access the firewall itself. The only access to the firewall is through a dedicated out-of-band management network which is neither accessible to the other networks on the firewall, nor is allowed to access anything other than the firewall itself. This segregates the management traffic from exploitable subnets greatly minimizes the firewall's vulnerability, even if an exploit is available.

Vulnerabilities can be software/hardware flaws or vulnerabilities of configuration / architecture. A proper configuration/architecture can help protect from both classes of vulnerabilities.

1.5.6 Management VPN Network

A back-end LAN-to-LAN VPN mesh will be installed using the Nokia Crypto Cluster VPN appliance. This device will serve to encrypt/decrypt all traffic between the various GIACE locations using IPSec standards. The VPN will initially traverse the open Internet, but can be migrated over to a dedicated Frame Relay network if security needs and reliability issues make it desirable. (A Frame Relay link typically is not as prone to the variable latency of the Internet, it is more difficult for someone to DoS your management connection, and it is usually considered more secure. Some may question the need for the VPN once Frame Relay is used. For me, the paranoia is justified for this connection since it services the company's crow jewels and because the tunnel is inexpensive and easy to maintain.)

The Nokia CC500 appliance will handle up to 6 Mbps of throughput using SHA-1 / 3DES and can be clustered with other devices to provide additive throughput, load balancing, and high-availability fail-over. Nokia also offers larger products in this line as the company grows.

1.5.7 Mirror Site / Disaster Recovery

A mirror site that duplicates the functionality of the primary site is important for several reasons. It supports the ability to maintain functionality if some natural (or manmade) catastrophe befalls the primary site. It allows for the possibility of globally load balancing access to the site and may allow for better response for end users. And it allows you to spend more money. (Seriously, the use of a mirror site must be balanced against its cost. Due diligence on the cost of any potential outages based upon time-line criteria, potential outage scenarios, and the possibility of each scenario must be undertaken to decide if a mirror site is warranted. It can be a powerful defense against outages but is not something to undertake lightly.)

At a minimum, all key servers, configurations, data, and etc. will be regularly backed up, periodically

checked for validity, stored in an off-site location, and (here's the key element) periodically used to restore during a mock disaster recovery exercises. These exercises can be time consuming and, frankly, a bit of a pain (try attending one every quarter). But the value is manifest because something ALWAYS goes wrong during the first attempt and the lessons learned make each successive attempt go smoother. Not ever practicing disaster recovery procedures for a system valued at \$200,000,000 per year is not the smartest version of employment roulette.

A mirror site may provide a built-in place to store the back-ups from the primary site as well as a place to test the disaster recovery procedures. With or without a mirror site, however, transfer of back-ups should always be done through a trusted agent that allows you to lock the storage units at your locations. Arcus Data Security is one company that can provide secure transfer, storage if needed, and other disaster recovery functions. They are a trusted name in the industry and widely used.

1.5.8 Corporate HQ / Offices

Corporate HQ will be deployed with a Checkpoint firewall to protect the T-1. The user subnets will be moved to private IP addresses and will be further divided by another firewall that will separate user subnets by function (Marketing, Finance, Operations, etc.). Split DNS will be extended to the corporate site to allow for an internal DNS server on a protected DMZ off the firewall. The external firewall will also be used to provide a separate mail DMZ that will supply mail relay and anti-virus checking for inbound and outbound mail.

No traffic will be allowed to be initiated inbound directly to the user subnets from the Internet. Outbound traffic from user subnets will be NAT-hidden and will be egress filtered to allow traffic to flow only on needed services (http, https, and others designated by the written security policy). Mail and DNS will be allowed from the user subnets to those designated DMZs.

The Nokia VPN appliance will be connected to the firewall via a crossover cable and the firewall will control access to the management VPN. Only designated user subnets will be allowed and only those services defined in the written security policy (SSH, SCP, etc.).

1.5.9 High-availability Components

The security architecture points out several locations that can benefit from high availability. These are not the only points where this can be accomplished, but are those points that should be investigated first when the architecture matures to that point. The border routers are an obvious first choice and the use of more than one ISP to protect against outages. The next point of high availability should be the primary firewalls. Checkpoint on Nokia IP hardware can accomplish a reliable standby firewall with the use of VRRP. The next level for redundancy is to provide fail-over capability on the PIX firewalls. There are other points of redundancy that can be figured in, including

usage of tandem switches for each subnet with redundant servers located on each switch, but this is not the main focus of this proposal. The first priority will be to implement a secure, flexible architecture and then to build upon that as needs grow. Single points of failure analysis can be performed later and high pay-off points can be address first.

© SANS Institute 2000 - 2002, Author retains full rights.

1.6 DEPLOYMENT PRIORITIES

Each section of the architecture has been explained with its purpose, hardware/software used, and overall policies of each. It was outlined earlier that the design was intended to be modular to allow for phased installation of the architecture. Following will be a discussion of the deployment priorities and reasons for each. This initial deployment plan and can be adjusted based upon management priorities, financial restrictions, unforeseen circumstances, etc.

Following are the recommendations for the initial deployment and the follow-on deployments as the security architecture matures. Obviously, the more that can be done initially up front will limit some of the costs involved overall. Staged deployments allow the finances to be spread out more, but the work effort involved and some costs are greater overall.

The initial deployment will include:

- Upgrading the border router IOS, configurations, and ACLs
- Installing the external (primary) firewall at the e-commerce location
- The VPN concentrator
- Partial Partner/Reseller and Customer Networks (web server and databases combined)
- Corporate HQ external firewall
- LAN-to-LAN management VPN
- Web and Database Management Network
- Security and Network Management Network

Recommendations to also include in the initial deployment if possible:

- Separation of the database servers to the outlined back-end networks with the protecting firewalls
- Split DNS
- Mail relay
- Netscreen firewall to protect the external DNS and mail relay servers
- Internal firewall at Corporate to separate functional user DMZs
- Off site storage / Disaster recovery vendor

Components to incorporate as the security architecture matures:

- High availability components
 - Secondary ISP and additional border router using HSRP
 - HA primary firewall
 - HA internal firewalls
 - Cluster added LAN-to-LAN devices to make VPN HA
- IDS
 - Border DMZ
 - Customer network
 - Partner/Reseller network
 - Database networks
 - Management Networks
- Captus for DoS and dDoS filtering
- Migrate VPN to Frame Relay

- Deployment of Mirror Site

© SANS Institute 2000 - 2002, Author retains full rights.

2.0 ASSIGNMENT 2 - SECURITY POLICY

2.1 ASSUMPTIONS

This section will outline the steps taken to configure and deploy components described in the preceding Security Architecture section of this paper. It will be assumed that the deployment priorities will be followed and that the design was not adjusted. The following sections will outline the specific deployment and configurations of the border router, the external (primary) firewall, and the LAN-to-LAN VPN appliances. Each will include an overview followed by details and specific ACLs/Security Policies.

2.2 BORDER ROUTER

2.2.1 Overview

The border router is a Cisco 7206 router that was previously owned by GIACE. It has a serial link that it uses to connect to their ISP and an Ethernet connection on the border DMZ. The configuration of this router involves upgrading its IOS version, installing ACLs to perform broad filtering functions, and some configuration to limit the vulnerabilities of the router itself.

2.2.2 Detailed steps to set ACLs & harden system

Upgrading the IOS and verifying interfaces

The border router was utilizing an older level of IOS code (The router's operating system) at version 11.1. The first step was to upgrade the IOS, so version 12.2.3 was installed. Once this was completed, the 'sh ver' command was used to verify. The pertinent portion is shown.

```
border-gw#sh ver
Cisco Internetwork Operating System Software
IOS (tm) C2600 Software (C2600-IS-M), Version 12.2(3), RELEASE SOFTWARE (fc2)
Copyright (c) 1986-2001 by cisco Systems, Inc.
```

Next, the interfaces were verified to be correct.

```
border-gw#show ip interface brief
```

Interface	IP-Address	OK?	Method	Status	Protocol
Ethernet0/0	208.a.1.1	YES	NVRAM	up	up
Serial0	unassigned	YES	unset	up	up
Serial0.1	208.a.2.1	YES	unset	up	up

Filtering for spoof-IPs

Next, access control lists were developed to filter out the illegitimate IP addresses recommended blocked by the SANS.org and to prevent spoofed use of internal IPs from external sources.

Internal interface anti-spoof

The following access list was created to prevent any traffic from leaving our border router that has an IP of a source other than from our public subnets (these are CIDR'd together as a single /27 for the ACL).

Creation of the ACL:

```
border-gw#no access-list 155
border-gw#access-list 155 permit ip 208.a.1.0 0.0.0.31 any
border-gw#access-list 155 deny ip any any log
```

ip-group 155 will be applied to Ethernet0/0 inbound. This has the effect of preventing any traffic that is not sourced from the border DMZ or the customer NAT subnet from being allowed outbound. This will prevent any compromised system on our networks from potentially spoofing their IP address to attack someone else on the Internet. It will also help prevent a mis-configuration/breach of a server/firewall from allowing private addresses outbound.

Applying the ACL:

```
border-gw#conf t
border-gw(config)#eth 0/0
border-gw(config-if)#ip-group 155 inbound
```

External interface anti-spoof

In the same vein, we built an access control list that will prevent any IP address that is inside the router from being allowed inbound from the Internet. This will help protect against some popular spoofing attacks that the attacker pretends to come from an IP address on the internal subnet (Smurf is a prime example).

Creation of the ACL:

```
border-gw#no access-list 165
border-gw#access-list 165 deny ip 208.a.1.0 0.0.0.31 any log
border-gw#access-list 165 permit ip any any
```

Applying the ACL:

```
border-gw#conf t
```

```
border-gw(config)#ser 0.1
border-gw(config-subif)#ip-group 165 inbound
```

Filtering Private and Reserved IP Addresses

Both access lists above will also include lines to filter out the private and reserved IPs (at SANS Institute recommendations). These lines will actually be part of the above ACLs but are separated here so that their functionality can be described more clearly.

(Both access-list 155 and 165 have the following lines appended to the top. The appropriate ACL number replaced the xxx.)

```
access-list xxx deny ip 0.0.0.0      0.255.255.255 any
access-list xxx deny ip 10.0.0.0     0.255.255.255 any
access-list xxx deny ip 127.0.0.0    0.255.255.255 any
access-list xxx deny ip 169.254.0.0  0.0.255.255 any
access-list xxx deny ip 172.16.0.0   0.15.255.255 any
access-list xxx deny ip 192.0.2.0    0.0.0.255 any
access-list xxx deny ip 192.168.0.0  0.0.255.255 any
access-list xxx deny ip 224.0.0.0    15.255.255.255 any
access-list xxx deny ip 240.0.0.0    7.255.255.255 any
access-list xxx deny ip 248.0.0.0    7.255.255.255 any
access-list xxx deny ip 255.255.255.255 0.0.0.0 any
access-list xxx permit ip any any
```

Other Router Settings

Some other configuration steps that will add additional filtering and to limit vulnerabilities of the router itself were taken. They are listed below.

To prevent the site from being used as a broadcast amplification point:

```
#no ip directed broadcast (applied to each interface)
```

This is default now on IOS versions 12.0 or higher.

Set strong passwords and set password encryption

```
#service password-encryption
```

! (Yes, I know it's easy to crack this. But that's why you still protect your config files ;)

© SANS Institute 2000 - 2002, Author retains full rights.

Force logon access from firewall only

Define access list

```
border-gw(config)#access-list 10 permit 208.a.1.2  
border-gw(config)#access-list 10 deny any
```

Apply access list

```
border-gw(config)#line vty 0 4  
border-gw(config)#access-class 10 in  
border-gw(config)#no exec  
border-gw(config)#password 7 0987654321FEDCBA  
border-gw(config)#login
```

Follow on steps in deployment of the security architecture will include a AAA server that will be set to authenticate for the router via the firewall.

The next step is to turn off any unneeded routing protocols and unneeded services that can give away unneeded information or allow too great of access. Disabling finger, ICMP unreachable messages, and CDP (Cisco Discovery Protocol) will limit information provided to an attacker. Disabling maintenance operation mode, ip source routing, multicast caching, small services, proxy ARP, IP redirects, and chargen will limit some other vulnerabilities.

```
border-gw(config)#no service finger  
border-gw(config)#no ip unreachable  
border-gw(config-if)#no cdp enable (on interfaces)  
border-gw(config)#no mop enabled  
border-gw(config)#no ip source-route  
border-gw(config)#no ip mroute-cache  
border-gw(config)#no service udp-small-servers  
border-gw(config)#no service tcp-small-servers  
border-gw(config)#no ip proxy-arp  
border-gw(config)#no ip redirects
```

2.3 EXTERNAL (PRIMARY) FIREWALL

2.3.1 Overview

The external (primary) firewall is a Nokia/Checkpoint firewall that protects the Customer and Partner/Reseller DMZs from the border and the Internet. The existing firewall was a Checkpoint firewall running on a Windows NT™ system. The license was reclaimed and used. The new external firewall was migrated to the UNIX-based Nokia IP330 appliance and utilizes five interfaces: an external interface on the border DMZ, an out-of-band management interface, an interface on the customer DMZ, an interface on the partner/reseller DMZ, and an interface crossed over into the VPN concentrator appliance.

Configuration of the firewall can be broken down into two distinct parts: configuring the Nokia appliance and configuring the Checkpoint firewall software.

2.3.2 Detailed steps used to configure Nokia appliance

The Nokia IP security appliance uses the IPSO operating system (OS). Nokia's hardened OS is based upon BSD UNIX and is designed to host the Checkpoint firewall software. The Nokia appliance was configured before the Checkpoint firewall with all necessary interface, routing, hostname information, etc. Nokia allows use of a Lynx browser running on localhost or via a web browser to their Voyager GUI to configure IPSO.

Connecting to the console interface of the Nokia IP330 and turning it on for the first time starts the configuration. It prompts for a hostname and some rudimentary IP information. The choice is given to use either a web browser (if you choose to enable it) or via the console port by entering 'lynx'. Lynx was used.

Lynx and Voyager are both menu driven, so describing the navigation in detail would be insulting. I will focus on the settings for this installation, the reasons, and some of the gotchas. A full configuration summary is provided in Appendix A.

Ensure proper code level

Since the Nokia is Unix-based, the uname command was used to verify the IPSO version.

```
External-FW[admin]# uname -a  
IPSO External-FW 3.3-FCS8 ericveum 651 05.14.2001-222300 i386
```

This box was loaded with IPSO 3.3-FCS8, which is fairly new and stable. IPSO 3.4.0 was

recently released but it does not have any critical fixes that are needed for this install and is too new to risk using.

If needed, migrating IPSO versions requires simply copying the package (via ftp, scp, etc.) onto the appliance and running the package update utilities.

Interface Configuration

Configuring the interfaces has a few steps and has a couple of gotchas. The first step is to turn on the needed interfaces and apply the changes. Forgetting to apply the changes before changing screens will cause the changes to be lost. Once the interfaces were enabled, each one was selected in turn and assigned the IP and description information. Once again, forgetting to apply changes can bite you. Under each interface, a physical interface-setting link leads to another menu. If not set, the Nokia will default to 10 Mbps /Half duplex. All interfaces were set to 100 Mbps and Full Duplex. Also buried in this section is the ability to shut off the logical interface. (If connectivity issues are driving you nuts, everything looks like it's turned on, but it still won't work, look here.)

A good practice is to double-check the configuration in the 'Show Configuration Summary' link. As mentioned, the full summary of the external (primary) firewall can be found in Appendix A. The interface configuration is shown here.

Physical int IP address	Speed	Duplex	Logical interface		Active / UP
=====	=====	=====	=====	=====	=====
eth-s2p1	100 Mbit	Full Duplex	External_Net	On Up	208.a.1.2/28
eth-s2p2	100 Mbit	Full Duplex	OOB_Mgmt_Net	On Up	10.1.100.1/24
eth-s3p1	100 Mbit	Full Duplex	PartnerReseller_Net	On Up	192.168.10.1/24
eth-s4p1	100 Mbit	Full Duplex	Customer_Net	On Up	192.168.20.1/24
eth-s5p1	100 Mbit	Full Duplex	VNP_Net	On Up	192.168.100.1/24
loop0			loop0c0	On Up	

Routing Configuration

The Nokia IP330 can be configured to use RIP or OSPF. This functionality was not used since the networks are small and easily maintained with static routes. This also eliminates one more potential for something to go wrong or to be exploited in some way.

Under 'Static Routes', the default gateway of 208.a.1.1 was entered.

Setting the Hostname and Host Address Assignments

In lynx, you can change your hostname. When changed, however, it will not show up on the command prompt until a reboot is performed. A crucial step is to ensure the Hostname is also represented in the Host Address Assignments with the IP address of the firewall interface that is licensed to Checkpoint. When installing the Checkpoint license key, it will perform a host look-up and will fail if a match for the IP by hostname is not found.

Under Host Address Assignment, External-FW was assigned an IP of 10.1.100.1 which matches the firewall's management interface and is the IP that the license was keyed to.

Setting the default filter

During the boot process, before the firewall daemon loads or if it fails, a filter can be designated. The external firewall was configured with a default filter to drop any traffic from any source for this purpose. This drop all filter is applied during all times the firewall daemon is not fully on line.

Installation of needed packages

These can be installed directly from an FTP server or copied onto the firewall using SCP or FTP and then installed. Checkpoint 4.1 SP3 and F-Secure SSH server and client 1.3.6.2 for IPSO were installed on the firewall.

Disable unneeded services

Even though the firewall rule base will not allow services such as telnet to the firewall (we will use SSH), the telnet daemon was turned off.

Likewise, any other unneeded services and routing protocols were disabled.

The SNMP write string was disabled and the read community was changed from the default of 'public'.

Lynx/Voyager Configuration Summary

See Appendix A

2.3.3 Detailed steps used to configure Checkpoint and set the security policy

Once IPSO is configured, the firewall software can be configured. Part of the IPSO configuration included installing the needed version of Checkpoint. To verify the correct version is installed, the 'fw ver' command was utilized.

```
External-FW[admin]# fw ver
```

This is Check Point VPN-1(TM) & FireWall-1(R) Version 4.1 Build 41821 [VPN + DES + STRONG]

Configuring and Licensing the Checkpoint Software

As shown above, we are running Checkpoint 4.1 SP3. Running the 'cpconfig' command performs some initial configuration steps and allows input of some base information concerning the management topology, license information, administration accounts, etc. The license was input, the user 'fwadmin' was created for administration of the firewall, and the firewall management was set to distributed.

Checkpoint breaks its components into three major pieces: firewall, management server, and configuration GUI. These can either be all on the same system, distributed fully, or partially distributed. This deployment has the management server located on the out of band management network at the address 10.1.100.200. This IP address was given to the firewall as its 'Master' and session keys were exchanged from each system to set up communication.

Select IPs was then designated in the management server as UI Clients and able to access it via the Policy Editor program.

Logging into the management server from one of the authorized IP addresses requires starting the Policy Editor software and providing valid a user name, password, and IP to the management server.



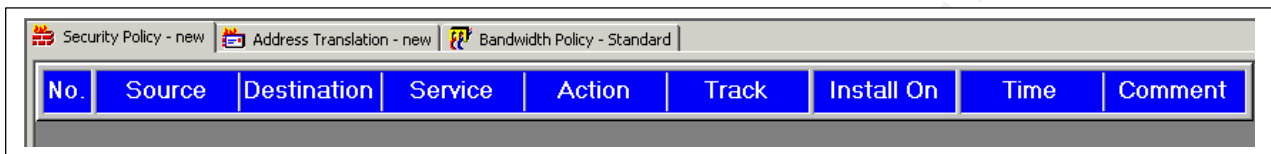
i
G
v

s
a
s
u
t

© SANS Institute 2000 - 2002, Author retains full rights.

The Checkpoint Security Policy Editor

Defining the security policy in Checkpoint is easiest if you follow some sort of organization. It is very easy for the rule base to get confusing, so using the comments column, consistent naming conventions, and color codes, etc. greatly increases its readability. The columns in the rule base are the Rule number, the Source, the Destination, Service, Action, Track, Install On, Time, and Comment.



No.	Source	Destination	Service	Action	Track	Install On	Time	Comment
-----	--------	-------------	---------	--------	-------	------------	------	---------

Rule Number: simply the number from 1 to however many rules in the policy numbered from top to bottom.

Source: The source IP address to match the rule against. Can be defined as an individual IP, a group of IPs, or a network.

Destination: The destination IP address to match the rule against. Defined identical to Source.

Service: Protocol and Port/Code/Etc. to match the rule against.

Action: Primarily used to Accept, Drop (silently), or Reject traffic. Also has different authentication options and encryption options used with VPNs.

Track: What level of logging to use. Ranges from None to Accounting.

Install On: A management server can control more than one firewall with the same rule base and specify which rules are installed on specific firewalls.

Time: Rules can be time based (only enabled or disabled during certain times).

Comment: Use. It will save you.

Configuration of the External (Primary) Firewall

Configuration of the primary firewall is broken down into the following steps.

- Removal of the default Implied Rules and settings
- Defining the objects (networks, hosts, groups) to be used in rule base

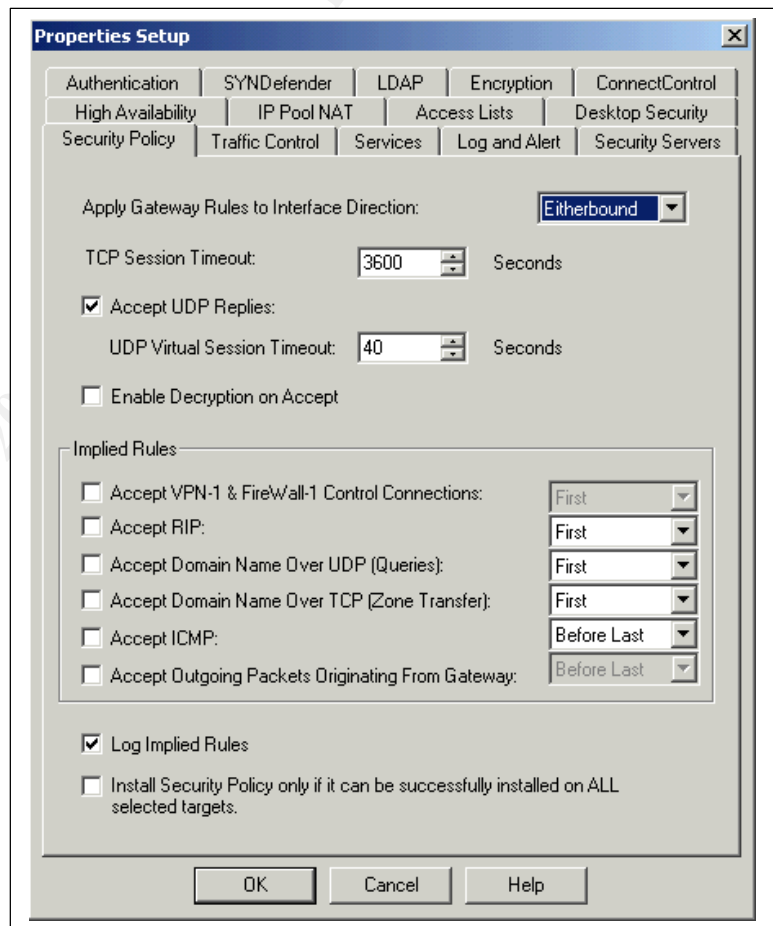
- Defining the firewall object and setting anti-spoof rules by interface
- Defining the rules to filter out invalid IPs
- Building the management rules
- Building the partner/reseller network rules
- Defining the NAT entries
- Building the customer network rules
- Building explicit rules to prevent customer and partner/reseller DMZ crossover
- Install rule base to firewall and verify

All of these steps are outlined below with explanations of why each step was undertaken. The next section will show the completed rule base.

Removal of the default Implied Rules and settings

Checkpoint, out of the box, has certain rules that are 'implied'. These are invisible unless you select View/Implied Rules from the menu bar. There are quite a few rules under the others created by these implied rules. These "helpful" little additions are too wide open for their own good and are best reconfigured from scratch so that they can be defined a little tighter.

To shut them off the Properties Setup/Security Policy tab was accessed via the menu policy/Properties. The selections seen in the figure show that the implied rules are now deselected.



Also, the checkbox for Log Implied Rules is now selected to ensure any drops from the system rules will show up in the log viewer. This saves no end of frustration because a firewall dropping something silently and not logging it can drive a troubleshooter a little batty.

The removal of the implied rules was double-checked by selecting the View/Implied Rules and observing that no additional rules appeared.

© SANS Institute 2000 - 2002, Author retains full rights.

Defining the objects (networks, hosts, groups) to be used in rule base

The following objects were defined for the firewall. Some were defined for use directly in the security policy tab. Others were for use behind the scenes in areas such as anti-spoof settings for the firewall object or NAT translation rules.

Customer Network Objects

Net_Cust_192.168.20.0	The private IP address of the public access network
Net_Cust_NAT_208.a.1.16	The public NAT IPs used for the VIPs.
Cust_WWW_VIP_192.168.20.17	The private IP address for the web server VIP.
Cust_SSL_VIP_192.168.20.18	The private IP address for the secure web server VIP.
Cust_Internal_DNS_192.168.20.50	The private IP address of the internal DNS server.
Cust_Servers_on_192.168.20.0	A group of the private IP addresses for the servers.
Cust_WWW_VIPNAT_208.a.1.17	The public NAT IP for the web server VIP.
Cust_SSL_VIPNAT_208.a.1.18	The public NAT IP for the secure web server VIP.

Partner/Reseller Network Objects

Net_PartResell_192.168.10.0	The private IP address of the partner/reseller network.
PartnerReseller_Srvr1 - Srvr9	Individual host objects for each server private IP address.
Partner_Servers	Group object of partner servers.
Reseller_Servers	Group object of reseller server.

Border Network & Internet Objects

Net_Border_208.a.1.0	The border DMZ subnet.
Border_gw	The border gateway router. (at IP 208.a.1.1)
NS_External_DNS_208.a.1.3	The IP address of the external DNS server. (behind the Netscreen firewall)
NS_SMTP_Relay_Srvr_208.a.1.4	The IP address of the external DNS server. (behind the Netscreen firewall)

VPN Network Objects

Net_VPN_192.168.100.0	The connecting network between the firewall and VPN.
VPN_concentrator	The private IP address of the VPN appliance connection to the firewall (at IP 192.168.100.2)

Management Network Objects

Net_OOB_mgmnt_10.1.100.0	The private IP address of the management subnet.
External-FW	The external firewall.
OOB_Management_Srvr	The firewall management server. (at IP 10.1.100.200)

Anti-Spoof Network Objects

(Although already filtered at the border router, these are defined on the firewall as well.)

invalid_0.0.0.0	Historical Broadcast
invalid_10.0.0.0	RFC 1918 Private Network
invalid_127.0.0.0	Loopback
invalid_169.254.0.0	Link Local Network
invalid_172.16.0.0	RFC 1918 Private Network
invalid_192.0.2.0	Test network
invalid_192.168.0.0	RFC 1918 Private Network
invalid_224.0.0.0	Class D Multicast
invalid_248.0.0.0	Class E Reserved
invalid_255.255.255.255	Broadcast
invalid_IPs	Group object that contains the invalid IP objects.

Allowed IPs for Interfaces

(the following group objects are used to set interface anti-spoofing in the firewall)


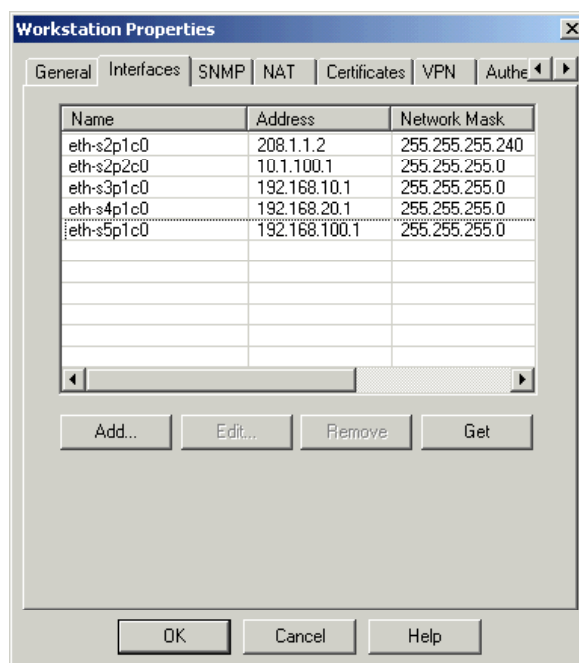
Allowed_int1	Group object defining networks allowed to interface 1.
Allowed_int2	Group object defining networks allowed to interface 2.
Allowed_int3	Group object defining networks allowed to interface 3.
Allowed_int4	Group object defining networks allowed to interface 4.
Allowed_int5	Group object defining networks allowed to interface 5.

Defining the firewall object and setting anti-spoof rules by interface

Once the needed network objects were created, the firewall object needed further definition and to set the anti-spoofing rules. This was accomplished by defining the firewall's interfaces by name, IP address, and subnet mask in the 'Interfaces' tab.

Selecting each interface and clicking Edit will enter the Interface Properties configuration. Under the Security tab, each interface has the Valid Addresses defined and has spoof-tracking set to log.

The purpose of this is to tell the firewall what IP addresses to expect as sources on that interface. This helps protect against an attacker using a spoofed address to pass through the firewall because the interface rules will only allow traffic from the expected sources. This is applied even before the Security Policy rules are checked.



The choices range from Any (no limitations) to Specific (which uses a defined network object). This net uses the address and mask of the Interface Properties tab to define the value (if there is a discrepancy between the actual physically assigned IP and mask and the one represented here, spoofing may drop some unexpected packets). The Others options sum up all the defined Valid Addresses on the rest of the firewall interfaces and allows all except them. Others is typically used on the Internet facing interface.









Each interface on the firewall is configured similarly to the example figure (this shows the management interface int2). The use of a network object allows valid

addresses to be added and removed easily by adding or subtracting from the 'Allowed' group.

The external interface was set to Others+ the Allowed_int1 group that includes the NAT IPs for the customer VIPs. Due to the way Checkpoint orders the way the system handles NAT, the NAT IPs can be 'seen' by the firewall on the internal interface. This causes the Others definition to exclude the NAT addresses so they need to be added in specifically using the Others+ option.








Defining the rules to filter out invalid IPs

In the security policy, the following rules were created to filter out unwanted source and destination IP addresses. This group can be easily added to in the future if you want to black-hole a particular IP or subnet. (For readability, this, and all subsequent depictions of rules will crop the Track, Install On, and Time Columns. These are set to Long, Gateways, and Any respectively.) The border router also filters out these invalid IPs outlined by the SANS Institute but the principle of defense in depth applies. It is very simple to apply some filtering here as well for an added measure of security. (The private 10 and 192.168 nets are not filtered on the firewall because they are valid subnets for the firewall's use, but the rest of the invalid IP list is filtered out. The border router filters the 10 and 192.168 nets.)

 invalid_IPs	 Any	 Any	 drop	Deny traffic from invalid source IP addresses
 Any	 invalid_IPs	 Any	 drop	Deny traffic to invalid source IP addresses

Building the management rules













Since the Implied Rules were removed, the following rules to allow management from the out-of-band management server (and only this server) were added. Checkpoint follows a top-down matching approach (Generally. There are some special cases with authentication, encryption and others.). Traffic from the management server to the firewall is matched and allowed. All other traffic to the firewall is dropped. As shown, the management server will be able to access the firewall on Checkpoint Firewall-1 management ports (TCP ports 256, 257, 258, and 259) and using SSH (TCP port 22).

 OOB_Management_Srvr	 External-FW	 FireWall1  ssh	 accept	Permit management from OOB network.
 Any	 External-FW  Net_OOB_mgmtnt_10.1.100.0	 Any	 drop	Stealth Rule Prevents traffic to firewall and to OOB Mgmt Net

An important note in the 'Stealth Rule' is the explicit drop of any traffic destined for the out-of-band management network as well as to the firewall itself.

Building the partner/reseller network rules










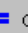


The rules for the partner and reseller network will only allow traffic from the VPN concentrator appliance via the crossover connection to the firewall. This will permit partners and resellers to securely VPN into the Cisco appliance and then access the servers over a secured link. The services available are http and https traffic to both sets of servers. The reseller servers also use a proprietary application on TCP port 9000 as previously mentioned. The partner servers have a similar application that uses TCP port 9001.

 VPN_concentrator	 Reseller_Servers	 tcp9000  http  https	 accept	Permit Resellers to access servers via VPN
 VPN_concentrator	 Partner_Servers	 tcp9001  http  https	 accept	Permit Partners to access servers via VPN

Defining the NAT entries


The customer servers are accessed from the Internet through the publicly NAT'd addresses of the load balancer VIPs. NAT rules are created in the Address Translation tab to translate the public IP address to the private IP address on the load balancer VIP behind the firewall.

This can be done using automatically created NAT entries by selecting NAT in the network objects. Like the implied rules, however, these are too generic and can be more succinctly defined by doing it manually. The following NAT entries match to specific destination and service before deciding to translate the destination.

Address Translation - External-FW						
No.	Original Packet			Translated Packet		
	Source	Destination	Service	Source	Destination	Service
1	 Any	 Cust_WWW_VIPNAT_208.a.1.17	 http	 Original	 Cust_WWW_VIP_192.168.20.17	 Original
2	 Any	 Cust_SSL_VIPNAT_208.a.1.18	 https	 Original	 Cust_SSL_VIP_192.168.20.18	 Original

Building the customer network rules.

The load balancer VIPs are allowed to be access from the Internet on the Security Policy tab by specific service. HTTP is allowed to the WWW VIP and HTTPS to the SSL VIP.









 Any	 Cust_WWW_VIPNAT_208.a.1.17	 http	 accept	Permit http traffic to Customer WWW VIP
 Any	 Cust_SSL_VIPNAT_208.a.1.18	 https	 accept	Permit https traffic to Customer SSL VIP

Since the customer servers provide an email confirmation of orders, traffic must be allowed through the firewall to support this. However, instead of just opening the floodgates and allowing the servers to blindly send outbound traffic anywhere and to query DNS from anything, the security architecture utilizes split DNS and mail relay. This allows the rule base to be defined very tightly.

The internal servers query the internal DNS server. This traffic does not need to traverse the firewall. The internal DNS server then utilizes the external DNS server on the border DMZ behind the Netscreen firewall. The rule to allow this traffic is defined between these two specific IPs (one of which is a private IP address that cannot traverse the border router's ACLs).

Once the servers know the destination IP for mail, they do not send the mail directly to the Internet. Instead, once again traffic is limited to the mail relay server on the border DMZ behind the Netscreen firewall.









All DNS and SMTP communication is outbound initiated from the behind the firewall and use private IPs that are filtered by the border router ACLs to prevent them from accessing the Internet directly (as well as the depicted firewall rules that also limit by the specific source and destination).

 Cust_Servers_on_192.168.20.0	 NS_SMTP_Relay_Srvr_208.a.1.4	 smtp	 accept	Permit Customer servers to send SMTP to relay server
 Cust_Internal_DNS_192.168.20.50	 NS_External_DNS_208.a.1.3	 dns	 accept	Permit Internal DNS to utilize the External DNS server

Building explicit rules to prevent customer and partner/reseller DMZ crossover

Toward the top of the rule base (before any accept rules) is placed rules to explicitly drop any traffic between the customer and partner/reseller DMZs. This may seem unnecessary if there are not any rules that allow traffic between the two, but it is a good safety tip to always explicitly drop any traffic that you absolutely do NOT want to occur. Checkpoint usage of the source 'Any' means just that: any. Rule bases tend to evolve and these explicit drops will ensure that an 'Any' inputted later does not do more than intended.

In fact, the rule that allows Any to the customer VIPs would allow the reseller/partner servers to match this as well. This may not be too great a problem initially, but it is an unnecessary door that can be nailed shut. Another explicit drop was the rule preventing any traffic to the out-of-band management network in the management rules.

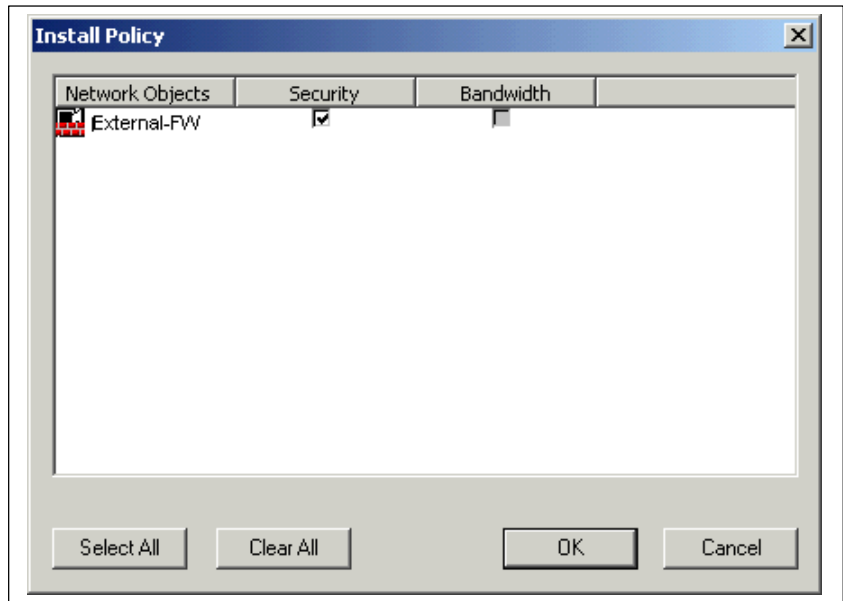
 Net_Cust_192.168.20.0	 Net_PartResell_192.168.10.0	 Any	 drop	Deny traffic from Customer DMZ to Partner/Reseller DMZ
 Net_PartResell_192.168.10.0	 Net_Cust_192.168.20.0	 Any	 drop	Deny traffic from Partner/Reseller DMZ to Customer DMZ

Install rule base to firewall and verify

The security policy will default to using the name Standard. This should be changed to something more meaningful. A standard naming convention will be used.

The initial security policy will be saved with the name External-FW.

Whenever a change is made to the rule base, the new policy will have its name appended with an eight digit number to designate the date. (i.e. External-FW09012001)



Once the rule base has been installed on the firewall, it is always a good idea to verify that it indeed made it. The easiest way (provided you use a naming convention on the security policy saves) is to console into the firewall and run the 'fw stat' command.

```
External-FW[admin]# fw stat
HOST    POLICY  DATE
localhost External-FW 11Sep2001 23:55:06 : [>eth-s2p2c0] [<eth-s2p2c0] [>eth-s2p2c0] [<eth-s2p2c0]
```

The firewall is running the policy named "External-FW" so everything is as it should be.

2.3.4 External Firewall Security Policy

Security Policy - External-FW					
No.	Source	Destination	Service	Action	Comment
1	Net_Cust_192.168.20.0	Net_PartResell_192.168.10.0	Any	drop	Prevents traffic from Customer Net to Partner/Reseller Net
2	Net_PartResell_192.168.10.0	Net_Cust_192.168.20.0	Any	drop	Prevents traffic from Partner/Reseller Net to Customer Net
3	invalid_IPs	Any	Any	drop	Deny traffic from invalid source IP addresses
4	Any	invalid_IPs	Any	drop	Deny traffic to invalid source IP addresses
5	OOB_Management_Svr	External-FW	FireWall1 ssh	accept	Permit management from OOB network.
6	Any	External-FW Net_OOB_mgmnt_10.1.100.0	Any	drop	Stealth Rule Prevents traffic to firewall and to OOB Mgmt Net
7	Any	Cust_WWWV_VIPNAT_208.a.1.17	http	accept	Permit http traffic to Customer WWW VIP
8	Any	Cust_SSL_VIPNAT_208.a.1.18	https	accept	Permit https traffic to Customer SSL VIP
9	VPN_concentrator	Reseller_Servers	tcp9000 http https	accept	Permit Resellers to access servers via VPN
10	VPN_concentrator	Partner_Servers	tcp9001 http https	accept	Permit Partners to access servers via VPN
11	Cust_Servers_on_192.168.20.0	NS_SMTP_Relay_Svr_208.a.1.4	smtp	accept	Permit Customer servers to send SMTP to relay server
12	Cust_Internal_DNS_192.168.20.50	NS_External_DNS_208.a.1.3	dns	accept	Permit Internal DNS to utilize the External DNS server
13	Any	Any	Any	drop	Clean-up Drop Rule

Address Translation - External-FW						
No.	Original Packet			Translated Packet		
	Source	Destination	Service	Source	Destination	Service
1	Any	Cust_WWWV_VIPNAT_208.a.1.17	http	Original	Cust_WWWV_VIP_192.168.20.17	Original
2	Any	Cust_SSL_VIPNAT_208.a.1.18	https	Original	Cust_SSL_VIP_192.168.20.18	Original

2.4 LAN-TO-LAN VPN APPLIANCES

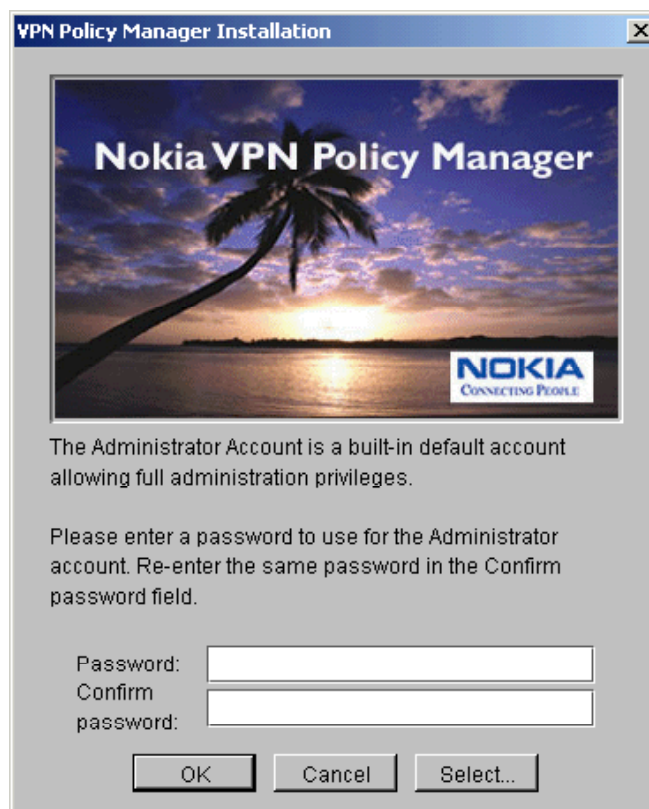
2.4.1 Overview

The LAN-to-LAN VPN network is deployed using hardware appliances from Nokia called the Crypto Cluster 500. The CC500 has two Ethernet interfaces and a console port. The configuration is basically limited to giving the device the needed IP information for its interfaces, a default gateway, a list of the IPs it encrypts for, and master station information. The devices use IPSec and default to using SHA-1/3DES and IKE. The setup is very simple and the devices are easy to maintain.

2.4.2 Steps used to configure and deploy the Nokia VPN appliances

The configuration for the Nokia CC is one by inputting the information into the management station which then allows you to copy/paste the setup into a console session on the device. Once the end device is configured, it will await connection from the management console using the stalled token (token time-out is configurable). The management station is configured to act as a Certificate Authority for the VPN endpoints and connects to them initially using the token. Communication between the management and endpoints is also encrypted. Additional endpoints can be configured and deployed at any time by the management station and then added to the existing VPN mesh.

The VPN is built using the 3.1(17) kernel on the Nokia CC endpoints and the v3.1 VPN policy manager.



The VPN tunnel was configured using the policy manager, pasting the configuration into the configuration wizard over a console session onto the devices, deploying and cabling the equipment, and pushing completion changes from the management station to the endpoints.

Configuration of the Corporate Nokia CC endpoint

The needed IP information is entered into the required blocks.

The inside IP address is the connection between the Corporate firewall and the Nokia CC. This is accomplished with the 10.200.1.0/30 subnet on a crossover connection. The corporate firewall has a static route configured to access the 10.1.0.0 / 16 network via the 10.200.1.1 address on the Nokia CC. The Nokia is configured with a static route to the 172.16.0.0 / 12 networks via the 10.200.1.2 IP address of the corporate firewall.

The outside IP address is the public address of the Nokia CC (from the ISP) and the Next Hop is the border router to the Internet.

The next step is to define the protected network(s). These are the subnets allowed to traverse the VPN from behind the Nokia CC. This is the locally defined encryption domain. There is no need to define the other end on this device. The other end of the tunnel is defined by entering a protected host group on the other Nokia CC. If three Nokia CCs are configured, each entering its own protected network creates a mesh of all three protected networks.

New Host Group

Description: CorpHQ - Protected

New Host Group Entry

☒ Enter IP address: 172.16.32.0

Subnet mask: 255.255.255.0

☐ Enter IP address range:

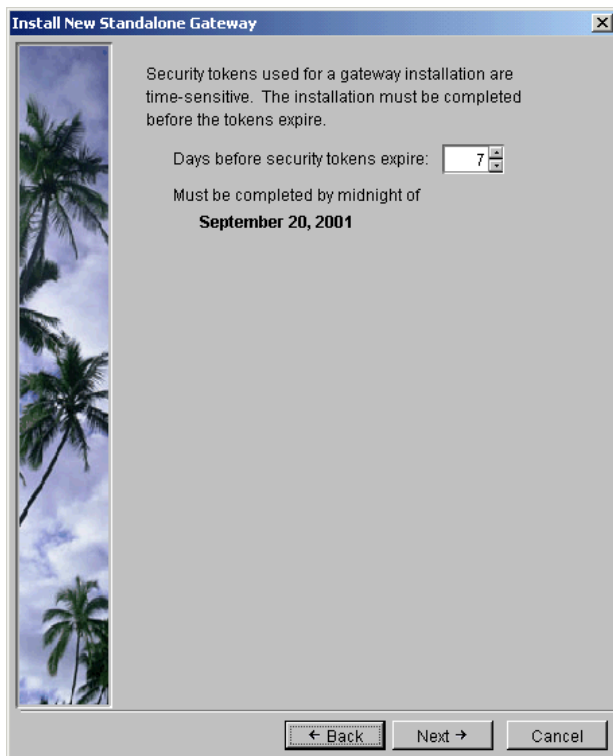
Starting IP address: 0.0.0.0

Ending IP address: 0.0.0.0

Number of IP addresses:

Comment: Corporate Operations / Engineering


OK Cancel



The length of time the token is available for connection is a configurable item. This defaults to seven days and was used.

This token is the pre-shared key that the management server uses to complete the installation when the Nokia CC endpoint is deployed to the remote location. If the token expires, the management station can generate a new token and then the new configuration must be installed on the endpoint before it will recognize the management server. If there is some doubt about whether the end-point will be installed and available to complete the installation within the 7 days, hedge higher. Having to coach someone over the phone to console into the device and then entering the configuration by hand is not enjoyable and isn't necessarily the most secure way to complete this process.

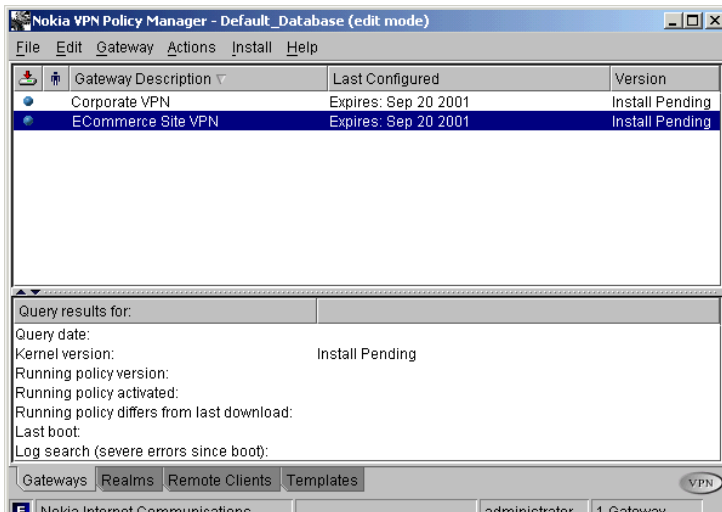
It is nice that the pre-shared key expires, but care needs to be taken to set the date to a reasonable expectation when deploying. For the deployment of the corporate end of the VPN tunnel, 7 days worked just fine. For the remote end at the e-commerce site, the days were set to 14 for deployment.



Once all the configuration information is entered, the management server creates the configuration for the VPN device. This can be cut/copied into the console session or entered by hand.

For this deployment, copy/pasting the configuration was used. In the Nokia CC console session, the command 'conf wiz' was entered. The prompt asked if we were sure? (y/N). We were sure, so we answered yes and the box rebooted into the configuration wizard. After pasting the generated configuration, the box was deployed to the field.

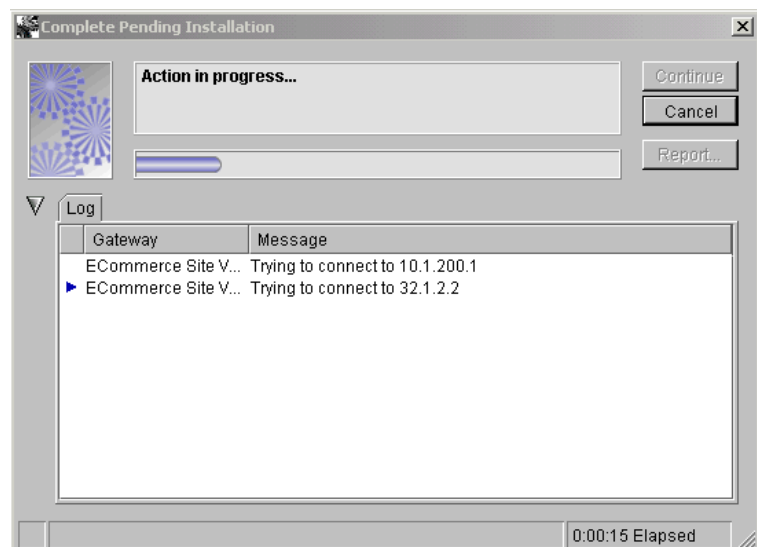
After both end-points were deployed, the policy manager shows them created with “Installation Pending”. Selecting Install / Complete Pending Installation completed the installation and the VPN was deployed.



The IPSec policy for the VPN tunnel is defined to use ESP with SHA-1 and DES.

Any traffic that does not match the filters for encryption are configured to be dropped by the endpoints.

The Nokia VPN appliances are configured to encrypt all traffic that meets the criteria of being in their encryption domain. The encryption domain is the Corporate perimeter / Engineering subnet at Corporate and the Web and Database and Security and Network management networks. Implementation of the protocols and ports will be handled by the firewalls.



w
t
o
d
d
O
s
W
S
M
L
a
t

3.0 ASSIGNMENT 3 – AUDIT OF SECURITY ARCHITECTURE

It is exceedingly important to conduct a thorough and ongoing audit of the security architecture to ensure that it is accomplishing what is expected of it. Simply configuring a firewall and putting it into production, for example, leaves open the possibility that some configuration steps were done incorrectly, the logic of the rule base was flawed in some way, the firewall doesn't perform the way the vendor claims (perish the thought), or something else was missed. It is important to take precautions to ensure the enforcement points are configured as securely as possible; and then to test them to confirm that you were successful.

The tool of choice for auditing is a network scanner. Scanners will attempt to connect to an IP address (or range of IP addresses) and determine what, if any, protocols and ports respond. Many scanners will then try to test known exploits against the discovered ports/protocols.

Conduct of a security audit, however, will be more than simply a network scan. For an audit to work best, it must be planned, conducted methodically, analyzed, and then used to correct the discovered vulnerabilities.

3.1 PLANNING THE SECURITY AUDITS

The Security Audit will be conducted in three parts: Pre-deployment, deployment, and ongoing. Pre-deployment audits will focus on auditing the individual components before they are put into the production environment. Deployment audits will concentrate on the interrelations of the components as part of the whole. They will also provide a baseline for use in corrective actions and as a comparison for the results of ongoing audits. Ongoing audits will be used to periodically take a snapshot of the security architecture and compare it to the baselines of previous audits, known changes/upgrades, and any corrective actions taken in between.

3.1.1 *Pre-deployment Audits*

Overview

Before installation into the security architecture, each component will be subjected to a pre-deployment audit. This audit will be conducted in three phases: An ACL/rule-base/configuration review, a scanning/sniffing in a lab environment, and needed corrective actions and rescanning. These steps will help ensure that the device is configured correctly and is a known quantity before it is installed into the larger architecture. This will help prevent the networks from having a great deal of vulnerabilities and unnecessary noise caused by unneeded services that can mask other problems.

An example of the pre-deployment audit will be demonstrated with the external (primary) firewall. The audit plan is laid out in the following paragraphs and is conducted in section 3.2.

External (Primary) Firewall Pre-Deployment Audit Plan

The external (primary) firewall will be audited prior to deployment. The steps for the pre-deployment audit will be:

- Configuration and Rule-base review

A second member of the security team will conduct the configuration and rule-base review. Provided documentation will be a copy of the network map, a list of the design requirements that shows the required IP and routing information, a copy of the written security policy, a printout of the interface summary and routing tables, a printout of the configuration summary, and a printout of the rule-base.

- Audit scanning and network sniffing in lab environment

The firewall will be connected to a scanning server on all of its interfaces. Each interface will then be subjected to a network scan to determine the ports available. Expected results will be compared to actual results. A sniffer will also be connected to the firewall's interfaces to determine if it is generating any traffic that is unwanted or unexpected. The firewall is also monitored on its interface dumps and the firewall logs to ensure it is seeing the traffic and handling it correctly during the scans.

- Correct discrepancies and rescan as needed

Any discrepancies discovered during the scanning and sniffing exercises will be analyzed and explained. Any corrections will be made and a rescan will be conducted if needed.

3.1.2 Deployment Audits

Overview

The focus of the deployment audit shifts from the individual component to the interrelation of the components as part of the whole. These audits will again test the security components, but will now attempt to test 'through' them as well. The networks themselves are also audited.

To conduct a proper audit, scans will be conducted from each possible network through to each other possible network. This can quickly get unwieldy if there are a lot of individual DMZs and multiple layers of networks and security enforcement points, so it is also important to prioritize the scans so that the more vulnerable possibilities are conducted first.

The Critical scans to be conducted

- Internet to border DMZ
- Internet/border to Customer public-access DMZ
- Internet/border to Partner/Reseller VPN-access DMZ
- Internet to Corporate
- Internet to all VPN appliances
- Local scan of Customer public-access DMZ
- Local scan of Partner/Reseller VPN-access DMZ
- Local scan of border DMZ

The Primary scans to be conducted.

- Customer public-access DMZ to Partner/Reseller VPN-access DMZ
- Partner/Reseller VPN -access DMZ to Customer public-access DMZ
- Cisco VPN DMZ to Customer and Partner/Reseller VPN-access DMZs
- All Internet, Customer, and Partner/Reseller DMZs to Security & Network Management DMZ (Re-confirm no access from any other DMZ.)
- Local scan of Security & Network Management DMZ
- Local scan of Web & Database Management DMZ
- Corporate private networks over the LAN-to-LAN VPN to the Management DMZs

The Secondary scans.

- Back-end subnets to database subnets and Web & Database Management network
- Web & Database Management to back-end subnets
- Local scan of remaining subnets

All scans will be collated into a report that specifies the time and date of the scan, the expected results, the output of the scan, the analysis of the audit, and any actions that will be undertaken. Security enforcement points that log will have their logs checked during the audit to ensure they are functioning correctly. IDS devices will also be checked to ensure proper functionality. The generated reports will be indexed by their placement in the above scan priority categories and numbered according to a matrix for quick reference. All categories of scans will be completed during the deployment audits. The categorization into critical, primary and secondary will be used to determine the overall severity for any discovered vulnerability and will also be used to determine the timing of ongoing audits.

3.1.3 Ongoing Audits

Ongoing audits will be conducted regularly and will be outlined in the written security policy. The policy will maintain who the authorized personnel are to conduct the audits, scheduling procedures, and the guidelines for reporting vulnerabilities for correction. The categories of critical, primary, and secondary will be used as a guide for scheduling. All critical scans will be conducted monthly. A portion of the primary scans will be conducted each month so that each is scanned at least once per quarter. A portion of the secondary scans will also be randomly checked each month so that each is scanned at least once annually. Most audits will be scheduled with prior notifications to administrators. At least one surprise audit will be conducted annually.

3.2 CONDUCTING THE SECURITY AUDITS

This section will outline some of the audits undertaken. The pre-deployment audit of the external (primary) firewall will be conducted as outline in the security audit planning section. The results of this audit will then be analyzed in the following sections. A representation of the deployment audits that test networks protected by the external (primary) firewall will also be represented and will also be analyzed in the following sections.

The audits will be conducted using TCPDUMP version 3.4 with libpcap 0.4. The network scans will be conducted using nmap V. 2.54BETA29 from www.insecure.org. Both tools will be run on a Linux server running 2.4.2-2 kernel.

3.2.1 Pre-deployment Audits

As outlined in section 3.1.1, the pre-deployment audits will be conducted on each security device individually before they are put into the security architecture. An example pre-deployment audit of the external (primary) firewall is provided. This is arguably the most critical security device to inspect and should be highly scrutinized. The audit will be conducted in three parts: configuration / rule-base review, audit scanning / sniffing, and analysis of audit and corrective actions.

Configuration and Rule-base review

A second member of the security team conducted the configuration and rule-base review. All the required documentation was presented. A discrepancy in the network diagram was discovered and corrected at this point (the Cisco 3000 VPN concentrator was erroneously shown to connect directly to the Partner/Reseller network). The corrected versions are what are represented in all the previous documentation. There was an error on the netmask setting for the management network in the firewall that was also corrected. The rule-base review resulted in no changes.

Audit scanning and network sniffing in lab environment

The audit outputs from the sniffs and scans are shown in this section. The analysis is handled in section 3.3.

The nmap scanner will use the flags to do stealth TCP and UDP scans (sS and sU) on all ports (-p 1-65535). The flags to prevent pinging (-P0) and reverse lookups (-n) are set and the operating system identification is used (-O).

The firewall logs and interfaces were inspected during all scans. An example of the logs is displayed for the scan of the management interface (the rest are left out because they are the same in the interest of brevity).

External (border) interface (208.a.1.1):

Network sniff –

```
#tcpdump -i eth0
Kernel filter, protocol ALL, TURBO mode (575 frames), datagram packet
socket
Tcpdump: listening on eth0

15:00:29:540106 I arp who-has 208.a.1.1 tell 208.a.1.2
15:01:10:540122 I arp who-has 208.a.1.1 tell 208.a.1.2
15:02:32:540098 I arp who-has 208.a.1.1 tell 208.a.1.2
15:03:13:540107 I arp who-has 208.a.1.1 tell 208.a.1.2
15:03:54:540131 I arp who-has 208.a.1.1 tell 208.a.1.2
      .      .      .
      .      .      .
      .      .      .
16:00:37:540109 I arp who-has 208.a.1.1 tell 208.a.1.2
16:01:18:540106 I arp who-has 208.a.1.1 tell 208.a.1.2
16:01:59:540118 I arp who-has 208.a.1.1 tell 208.a.1.2
16:02:40:540109 I arp who-has 208.a.1.1 tell 208.a.1.2
16:03:21:540105 I arp who-has 208.a.1.1 tell 208.a.1.2

64 packets received by filter
0 packets dropped by kernel
```

Nmap scan –

TCP PORTS

```
#nmap -sS -p 1-65535 -O -P0 -n 208.a.1.2

Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )
Warning: OS detection will be MUCH less reliable because we did not
find at least 1 open and 1 closed TCP port
All 65535 scanned ports on (208.a.1.2) are: filtered
Too many fingerprints match this host for me to give an accurate OS
guess
```

Nmap run completed - 1 IP address (1 host up) scanned in 32960 seconds

© SANS Institute 2000 - 2002, Author retains full rights.

UDP PORTS

```
#nmap -sU -p 1-65535 -O -P0 -n 208.a.1.2
```

```
Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )  
Warning: OS detection will be MUCH less reliable because we did not  
find at least 1 open and 1 closed TCP port  
All 65535 scanned ports on (208.a.1.2) are: filtered  
Too many fingerprints match this host for me to give an accurate OS  
guess
```

```
Nmap run completed - 1 IP address (1 host up) scanned in 15043 seconds
```

Customer network interface (192.168.20.1):

Network sniff –

```
#tcpdump -i eth0  
Kernel filter, protocol ALL, TURBO mode (575 frames), datagram packet  
socket  
Tcpdump: listening on eth0  
  
0 packets received by filter
```

Nmap scan –

TCP PORTS

```
#nmap -sS -p 1-65535 -O -P0 -n 192.168.20.1
```

```
Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )  
Warning: OS detection will be MUCH less reliable because we did not  
find at least 1 open and 1 closed TCP port  
All 65535 scanned ports on (192.168.20.1) are: filtered  
Too many fingerprints match this host for me to give an accurate OS  
guess
```

```
Nmap run completed - 1 IP address (1 host up) scanned in 29956 seconds
```

UDP PORTS

```
#nmap -sU -p 1-65535 -O -P0 -n 192.168.20.1
```

```
Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )  
Warning: OS detection will be MUCH less reliable because we did not  
find at least 1 open and 1 closed TCP port  
All 65535 scanned ports on (192.168.20.1) are: filtered  
Too many fingerprints match this host for me to give an accurate OS  
guess
```

Nmap run completed - 1 IP address (1 host up) scanned in 14432 seconds

© SANS Institute 2000 - 2002, Author retains full rights.

Partner/Reseller interface (192.168.10.1):

Network sniff –

```
#tcpdump -i eth0
Kernel filter, protocol ALL, TURBO mode (575 frames), datagram packet
socket
Tcpdump: listening on eth0

0 packets received by filter
```

Nmap scan –

TCP PORTS

```
#nmap -sS -p 1-65535 -O -P0 -n 192.168.10.1

Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )
Warning: OS detection will be MUCH less reliable because we did not
find at least 1 open and 1 closed TCP port
All 65535 scanned ports on (192.168.10.1) are: filtered
Too many fingerprints match this host for me to give an accurate OS
guess

Nmap run completed - 1 IP address (1 host up) scanned in 31654 seconds
```

UDP PORTS

```
#nmap -sU -p 1-65535 -O -P0 -n 192.168.10.1

Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )
Warning: OS detection will be MUCH less reliable because we did not
find at least 1 open and 1 closed TCP port
All 65535 scanned ports on (192.168.10.1) are: filtered
Too many fingerprints match this host for me to give an accurate OS
guess

Nmap run completed - 1 IP address (1 host up) scanned in 11769 seconds
```

VPN interface (192.168.100.1):

Network sniff –

```
#tcpdump -i eth0
Kernel filter, protocol ALL, TURBO mode (575 frames), datagram packet
socket
Tcpdump: listening on eth0

0 packets received by filter
```

Nmap scan –

TCP PORTS

```
#nmap -sS -p 1-65535 -O -P0 -n 192.168.100.1

Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )
Warning: OS detection will be MUCH less reliable because we did not
find at least 1 open and 1 closed TCP port
All 65535 scanned ports on (192.168.100.1) are: filtered
Too many fingerprints match this host for me to give an accurate OS
guess

Nmap run completed - 1 IP address (1 host up) scanned in 31004 seconds
```

UDP PORTS

```
#nmap -sU -p 1-65535 -O -P0 -n 192.168.100.1

Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )
Warning: OS detection will be MUCH less reliable because we did not
find at least 1 open and 1 closed TCP port
All 65535 scanned ports on (192.168.100.1) are: filtered
Too many fingerprints match this host for me to give an accurate OS
guess

Nmap run completed - 1 IP address (1 host up) scanned in 14797 seconds
```


Management interface (10.1.100.1):

Network sniff –

```
#tcpdump -i eth0
Kernel filter, protocol ALL, TURBO mode (575 frames), datagram packet
socket
Tcpdump: listening on eth0

0 packets received by filter
```

Nmap scan –

The scan of the management interface will be conducted from two source IPs. One that is on the management network but not allowed to access the firewall and also from the management server's IP address.

Scan from an IP on the management subnet other than the management server.
(10.1.100.100)

TCP PORTS

```
#nmap -sS -p 1-65535 -O -P0 -T insane -n 10.1.100.1

Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )
Warning: OS detection will be MUCH less reliable because we did not
find at least 1 open and 1 closed TCP port
All 65535 scanned ports on (192.168.100.1) are: filtered
Too many fingerprints match this host for me to give an accurate OS
guess

Nmap run completed - 1 IP address (1 host up) scanned in 3400 seconds
```

UDP PORTS

```
#nmap -sU -p 1-65535 -O -P0 -T insane -n 10.1.100.1

Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )
Warning: OS detection will be MUCH less reliable because we did not
find at least 1 open and 1 closed TCP port
All 65535 scanned ports on (192.168.100.1) are: filtered
Too many fingerprints match this host for me to give an accurate OS
guess

Nmap run completed - 1 IP address (1 host up) scanned in 3290 seconds
```

VIEW OF FIREWALL LOGS OF SCAN

No.	Time	Inter.	Origin	Type	Action	Service	Source	Destination	Prot..	Rule	S_Port
267544	17:34:29	eth-s2p2c0	External-FW	log	drop	nameserver	10.1.100.100	External-FW	tcp	6	52076
267545	17:34:29	eth-s2p2c0	External-FW	log	drop	8	10.1.100.100	External-FW	tcp	6	52076
267546	17:34:29	eth-s2p2c0	External-FW	log	drop	ni-ftp	10.1.100.100	External-FW	tcp	6	52076
267547	17:34:29	eth-s2p2c0	External-FW	log	drop	12	10.1.100.100	External-FW	tcp	6	52076
267548	17:34:29	eth-s2p2c0	External-FW	log	drop	TACACSplus	10.1.100.100	External-FW	tcp	6	52077
267549	17:34:29	eth-s2p2c0	External-FW	log	drop	mpm	10.1.100.100	External-FW	tcp	6	52077
267550	17:34:29	eth-s2p2c0	External-FW	log	drop	10	10.1.100.100	External-FW	tcp	6	52077
267551	17:34:29	eth-s2p2c0	External-FW	log	drop	6	10.1.100.100	External-FW	tcp	6	52077
267552	17:34:29	eth-s2p2c0	External-FW	log	drop	28	10.1.100.100	External-FW	tcp	6	52077
267553	17:34:29	eth-s2p2c0	External-FW	log	drop	msh	10.1.100.100	External-FW	tcp	6	52077
267554	17:34:29	eth-s2p2c0	External-FW	log	drop	nameserver	10.1.100.100	External-FW	tcp	6	52077
267555	17:34:29	eth-s2p2c0	External-FW	log	drop	8	10.1.100.100	External-FW	tcp	6	52077
267556	17:34:29	eth-s2p2c0	External-FW	log	drop	ni-ftp	10.1.100.100	External-FW	tcp	6	52077
267557	17:34:29	eth-s2p2c0	External-FW	log	drop	12	10.1.100.100	External-FW	tcp	6	52077
267558	17:34:29	eth-s2p2c0	External-FW	log	drop	37157	10.1.100.100	External-FW	tcp	6	52084

Scan from the IP of the management server.
(10.1.100.200)

TCP PORTS

```
[root@localhost bin]# nmap -sS -p 1-65535 -O -P0 -T insane -n 10.1.100.1
```

```
Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )
Interesting ports on (10.1.100.1):
(The 65512 ports scanned but not shown below are in state: closed)
Port      State      Service
22/tcp    open      ssh
256/tcp   open      rap
257/tcp   open      set
258/tcp   open      yak-chat
259/tcp   open      esro-gen
261/tcp   open      nsiiops
262/tcp   open      arcisdms
264/tcp   open      bgmp
265/tcp   open      unknown
443/tcp   open      https
900/tcp   open      unknown
1024/tcp  open      kdm
1025/tcp  open      listen
1026/tcp  open      nterm
1027/tcp  open      unknown
1028/tcp  open      unknown
1029/tcp  open      unknown
18183/tcp open      unknown
18184/tcp open      unknown
18185/tcp open      unknown
```

```
18207/tcp  open          unknown
19190/tcp  open          unknown
19191/tcp  open          unknown
```

Remote OS guesses: Check Point FireWall-1 4.0 SP-5 (IPSO build), Nokia IPSO 3.2-3.2.1 releng 783-849, NOKIA IPSO 3.2 Running Checkpoint Firewall-1, Nokia IPSO 3.2-fcs4 releng 783 (FreeBSD Based)
Uptime 1.910 days (since Mon Sep 17 01:12:31 2001)

Nmap run completed -- 1 IP address (1 host up) scanned in 12098 seconds

(*** These are more than was expected, see analysis section. ***)

UDP PORTS

```
#nmap -sU -p 1-65535 -O -P0 -T insane -n 10.1.100.1
```

Starting nmap V. 2.54BETA29 (www.insecure.org/nmap/)
Warning: OS detection will be MUCH less reliable because we did not find at least 1 open and 1 closed TCP port
All 65535 scanned ports on (192.168.100.1) are: filtered
Too many fingerprints match this host for me to give an accurate OS guess

Nmap run completed - 1 IP address (1 host up) scanned in 3290 seconds

Analysis of audits, corrections of discrepancies, and rescans if needed

Once all the audit data is compiled, it must be reviewed and analyzed to ensure the firewall is providing the proper function. Any discrepancies must be researched and corrected. Of the scans and sniff outputs, the scan of the management interface and the sniff of the external interface provided the most interesting output. The full analysis and corrective actions are detailed in the analysis section 3.3.

3.2.2 Deployment Audits

The deployment audits outlined in section 3.1.2 test the security of the overall architecture by testing the subnets and the traffic flows through the various security enforcement points. As the audit plan specified, the audits are grouped by category by the priority of the scan. The critical group has priority over the primary group and the primary has priority over the secondary group. A subset of the scans is included. The analysis of the scans is presented in section 3.3.

Critical Scans (example subset)

Internet/border to Customer public-access DMZ & Internet/border to Partner/Reseller VPN-access DMZ

- Internet/border to Customer public-access DMZ

Internet/border to Public NAT IP subnet (208.a.1.16/28)

```
[root@localhost bin]# nmap -sS -p 1-65535 -P0 -n 208.a.1.16/28

Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )
All 65535 ports on (208.1.1.16) are: filtered
Interesting ports on (208.1.1.17):
(The 65534 ports scanned but not shown below are in state: filtered)
Port      State      Service
80/tcp    open      http
Interesting ports on (208.1.1.18):
(The 65534 ports scanned but not shown below are in state: filtered)
Port      State      Service
443/tcp   open      https
All 65535 ports on (208.1.1.19) are: filtered
All 65535 ports on (208.1.1.20) are: filtered
All 65535 ports on (208.1.1.21) are: filtered
All 65535 ports on (208.1.1.22) are: filtered
All 65535 ports on (208.1.1.23) are: filtered
All 65535 ports on (208.1.1.24) are: filtered
All 65535 ports on (208.1.1.25) are: filtered
All 65535 ports on (208.1.1.26) are: filtered
All 65535 ports on (208.1.1.27) are: filtered
All 65535 ports on (208.1.1.28) are: filtered
All 65535 ports on (208.1.1.29) are: filtered
All 65535 ports on (208.1.1.30) are: filtered
All 65535 ports on (208.1.1.31) are: filtered

Nmap run completed -- 16 IP addresses (16 hosts up) scanned in 23098
seconds
```

Internet/border to Private IP subnet (192.168.20.0/24)

```
[root@localhost bin]# nmap -sS -p 1-65535 -P0 -n 192.168.20.0/24
```

```
Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )
```

```
All 65535 ports on (192.168.20.0) are: filtered
```

```
All 65535 ports on (192.168.20.1) are: filtered
```

```
All 65535 ports on (192.168.20.2) are: filtered
```

```
All 65535 ports on (192.168.20.3) are: filtered
```

```
All 65535 ports on (192.168.20.4) are: filtered
```

```
. . .  
. . .
```

```
(Summarized to save trees. All were filtered.)
```

```
. . .  
. . .
```

```
All 65535 ports on (192.168.20.253) are: filtered
```

```
All 65535 ports on (192.168.20.254) are: filtered
```

```
All 65535 ports on (192.168.20.255) are: filtered
```

```
Nmap run completed -- 256 IP addresses (256 hosts up) scanned in 25098  
seconds
```

- Internet/border to Partner/Reseller VPN-access DMZ

Border to Private IP subnet (192.168.10.0/24)

```
[root@localhost bin]# nmap -sS -p 1-65535 -P0 -n 192.168.10.0/24
```

```
Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )
```

```
All 65535 ports on (192.168.20.0) are: filtered
```

```
All 65535 ports on (192.168.20.1) are: filtered
```

```
All 65535 ports on (192.168.20.2) are: filtered
```

```
All 65535 ports on (192.168.20.3) are: filtered
```

```
All 65535 ports on (192.168.20.4) are: filtered
```

```
. . .  
. . .
```

```
(Summarized to save trees. All in between were filtered.)
```

```
. . .  
. . .
```

```
All 65535 ports on (192.168.20.253) are: filtered
```

```
All 65535 ports on (192.168.20.254) are: filtered
```

```
All 65535 ports on (192.168.20.255) are: filtered
```

```
Nmap run completed -- 256 IP addresses (256 hosts up) scanned in 24907  
seconds
```

Primary Scans (example subset)

- Cisco VPN DMZ to Customer and Partner/Reseller VPN-access DMZs

Cisco VPN device (192.168.100.2) to Partner/Reseller VPN-access DMZ (192.168.10.0/24)

```
[root@localhost bin]# nmap -sS -p 1-65535 -P0 -n 192.168.10.0/24
```

```
Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )
All 65535 ports on (192.168.10.0) are: filtered
Interesting ports on (192.168.10.1):
(The 65532 ports scanned but not shown below are in state: filtered)
Port      State      Service
80/tcp    open       http
443/tcp   open       https
9000/tcp  open       unknown
Interesting ports on (192.168.10.2):
(The 65532 ports scanned but not shown below are in state: filtered)
Port      State      Service
80/tcp    open       http
443/tcp   open       https
9000/tcp  open       unknown
Interesting ports on (192.168.10.3):
(The 65532 ports scanned but not shown below are in state: filtered)
Port      State      Service
80/tcp    open       http
443/tcp   open       https
9000/tcp  open       unknown
Interesting ports on (192.168.10.4):
(The 65532 ports scanned but not shown below are in state: filtered)
Port      State      Service
80/tcp    open       http
443/tcp   open       https
9000/tcp  open       unknown
Interesting ports on (192.168.10.5):
(The 65532 ports scanned but not shown below are in state: filtered)
Port      State      Service
80/tcp    open       http
443/tcp   open       https
9001/tcp  open       unknown
Interesting ports on (192.168.10.6):
(The 65532 ports scanned but not shown below are in state: filtered)
Port      State      Service
80/tcp    open       http
443/tcp   open       https
9001/tcp  open       unknown
Interesting ports on (192.168.10.7):
(The 65532 ports scanned but not shown below are in state: filtered)
Port      State      Service
80/tcp    open       http
443/tcp   open       https
9001/tcp  open       unknown

Interesting ports on (192.168.10.8):
```

```

(The 65532 ports scanned but not shown below are in state: filtered)
Port      State      Service
80/tcp    open       http
443/tcp   open       https
9000/tcp  open       unknown
Interesting ports on (192.168.10.9):
(The 65532 ports scanned but not shown below are in state: filtered)
Port      State      Service
80/tcp    open       http
443/tcp   open       https
9001/tcp  open       unknown
All 65535 ports on (192.168.10.10) are: filtered
All 65535 ports on (192.168.10.11) are: filtered
All 65535 ports on (192.168.10.12) are: filtered
.
.
.
(Summarized to save trees. All in between were filtered.)
.
.
.
All 65535 ports on (192.168.10.253) are: filtered
All 65535 ports on (192.168.10.254) are: filtered
All 65535 ports on (192.168.10.255) are: filtered

Nmap run completed -- 256 IP addresses (256 hosts up) scanned in 31098
seconds

```

3.2.3 Ongoing Audits

As detailed in section 3.1.3, the ongoing audits will be conducted at intervals. These will be the same as the deployment audits and will be date stamped, correlated with input from any changes between scans, and used to provide input into any needed corrective actions.

3.3 ANALYZING THE SECURITY AUDITS

3.3.1 Pre-deployment Audit

External (Primary) Firewall Pre-Deployment Audit Analysis

The audit of the firewall was conducted with full TCP and UDP scans on each interface in turn. Additionally, each interface was sniffed and the firewall logs were reviewed to ensure the traffic was indeed reaching it. The outputs were presented in section 3.2.1.

Most of the scans / sniffs revealed what was expected. The sniff on the external interface showed that the firewall was ARPing periodically. This is normal and expected for the firewall to be looking for its default gateway. This will stop when it is put into production and

receives a reply from the border router.

The full scans against the firewall interface reported that all ports were filtered (This means that nmap received no response back. Nmap will report 'closed' if it receives a reset.). This is as it should be and the firewall gives away no unnecessary information on any of its protected or external interfaces.

The exception, of course, was when the firewall was scanned from the management server IP address. This scan was expected to see the ports allowed in the management rule of the rule-base (rule #5). These ports were TCP ports 256, 257, 258, and 259 which are used by Checkpoint. Instead, several additional ports were discovered open.

Referring to the Checkpoint courseware, documentation, and www.phoneboy.com (an excellent Checkpoint FAQ site) gathered information on the discovered ports that were noticed. Checkpoint apparently includes filters that allow some traffic to the firewall from the management server by default that is not in the rule-base. These would be implied, implied rules that cannot be removed in the rulebase.

Entering lynx and ensuring that Voyager Web Access is set to be disabled can remove the visibility of ports 80 and 443. (This will also stop localhost access, but can be re-enabled for each use at the command line with the 'voyager' and then 'lynx' commands.)

The other ports are used for a variety of functions (mostly VPN or client authentication) that are not in use by this firewall and can be shut down by editing some files in the \$FWDIR/conf directory.

Editing of the fwauthd.conf file.

```
#vi $FWDIR/conf/fwauthd.conf
```

```
;21          in.aftpd          wait      0
;80          in.ahttpd         wait      0
;513         in.arlogind       wait      0
;25          in.asmtpd         wait      0
;23          in.atelnetd       wait      0
259          in.aclientd       wait      259
;10081       in.lhttpd         wait      0
;900         in.ahclientd      wait      900
0            in.pingd          respawn 0
0            polsrvd           respawn 0
0            netsod            respawn 0
```

```
~
~
~
```

```
fwauthd.conf: unmodified: line 1.
```


Editing of the fwopsec.conf file.

```
#vi $FWDIR/conf/fwopsec.conf
```

```
;sam_server      auth_port      18183
;lea_server      auth_port      18184
;ela_proxy       auth_port      18187
;ela_proxy       auth_type      ssl_opsec
ela_proxy        fwd_machine    localhost

# authenticated connections for servers
# server          <server IP>      <service port>      auth_opsec

;server          127.0.0.1          18181                auth_opsec
;server          127.0.0.1          18182                auth_opsec

sam_allow_remote_requests no

;netso_server     auth_port      19191
;netso_server     auth_type      ssl_opsec

~
~
~
fwopsec.conf: unmodified: line 1.
```

After all the corrections were made, a rescan of the interfaces for the discovered ports was conducted...

```
[root@localhost bin]# nmap -sS -p 200-500 -O -P0 -T insane -n 10.1.100.1

Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )
Interesting ports on (10.1.100.1):
(The 297 ports scanned but not shown below are in state: filtered)
Port      State      Service
256/tcp    open       rap
257/tcp    open       set
258/tcp    open       yak-chat
259/tcp    open       esro-gen

Remote OS guesses: Check Point FireWall-1 4.0 SP-5 (IPSO build), Nokia
IPSO 3.2-3.2.1 releng 783-849, NOKIA IPSO 3.2 Running Checkpoint
Firewall-1, Nokia IPSO 3.2-fcs4 releng 783 (FreeBSD Based)
Uptime 0.011 days (since Wed Sep 19 02:56:31 2001)

Nmap run completed -- 1 IP address (1 host up) scanned in 119 seconds
```

```
[root@localhost bin]# nmap -sS -p 900-1050 -O -P0 -T insane -n
10.1.100.1

Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )
Warning: OS detection will be MUCH less reliable because we did not
find at least 1 open and 1 closed TCP port
Interesting ports on (10.1.100.1):
(The 150 ports scanned but not shown below are in state: filtered)
Port      State      Service
900/tcp    closed     unknown

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

Nmap run completed -- 1 IP address (1 host up) scanned in 37 seconds


[root@localhost bin]# nmap -sS -p 18000-18210 -P0 -T insane -n
10.1.100.1

Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )
All 211 scanned ports on (10.1.100.1) are: filtered

Nmap run completed -- 1 IP address (1 host up) scanned in 31 seconds


[root@localhost bin]# nmap -sS -p 19190-19191 -P0 -T insane -n
10.1.100.1

Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap/ )
Interesting ports on (10.1.100.1):
Port      State      Service
19190/tcp  filtered   unknown
19191/tcp  filtered   unknown

Nmap run completed -- 1 IP address (1 host up) scanned in 2 seconds
```

Review of rescan

Now, only the expected and desired ports are available, and only from the management server. Audit completed.

3.3.2 Deployment Audits

Critical Scans (example subset)

The subset critical scans were from a lab mock-up Internet and from the border router IP. The subnets scanned were the public IP addresses used for NAT'ing of the Customer DMZ, the private IP addresses of the Customer DMZ, and the Partner/Reseller DMZ. The output of the scans were presented in section 3.2.2.

The scan of the public NAT IPs originally did not return the ports expected. The rule-base was set to allow web traffic to the private IP addresses of the VIPs instead of the public NAT IP address. Since Checkpoint matched the traffic to the rule-base before applying the NAT, the traffic did not match the intended rule and was dropped. This was noticed in the rule-base, corrected, and rescanned. Observing the firewall logs, network sniffing, and scan results now confirmed that the desired IP addresses of .17 and .18 responded to the desired ports and no other responses were received by the scanner. Results match desired status.

The scans of the private IP'd subnets were filtered by the border router from the mock-up Internet IPs and showed up on the scanner as 'filtered'. Scanning from the border subnet revealed the same results. The results were those desired.

Primary Scans (example subset)

A scan from the internal Cisco VPN IP address was conducted to test the access provided to the Partner/Reseller network through the client VPN links. This scan revealed the ports for web traffic and the proprietary ports for the Partner and the Reseller servers. The results were those desired so no further action was required.

3.4 CORRECTING VULNERABILITIES DISCOVERED DURING THE SECURITY AUDITS

Simply conducting a security audit is not enough by itself. The output needs to be analyzed to determine what was actually discovered. Once the analysis is complete, corrective actions need to be undertaken. Without this step, the audit does not accomplish much more than to confirm vulnerabilities. As simple as it sounds, correcting what was discovered is often one of the hardest and most time consuming steps. This is especially true of local scans of server subnets where the devices are not under the direct control of the security team. This is usually an exercise of technical research, diplomacy, scheduling, and maybe even some cajolery. The corrective actions needed for the audit examples conducted in section 3.2 were included as part of the steps in the analysis section of 3.3. This section is simply explaining the critical need for this step and to caution that fixes aren't always simple.

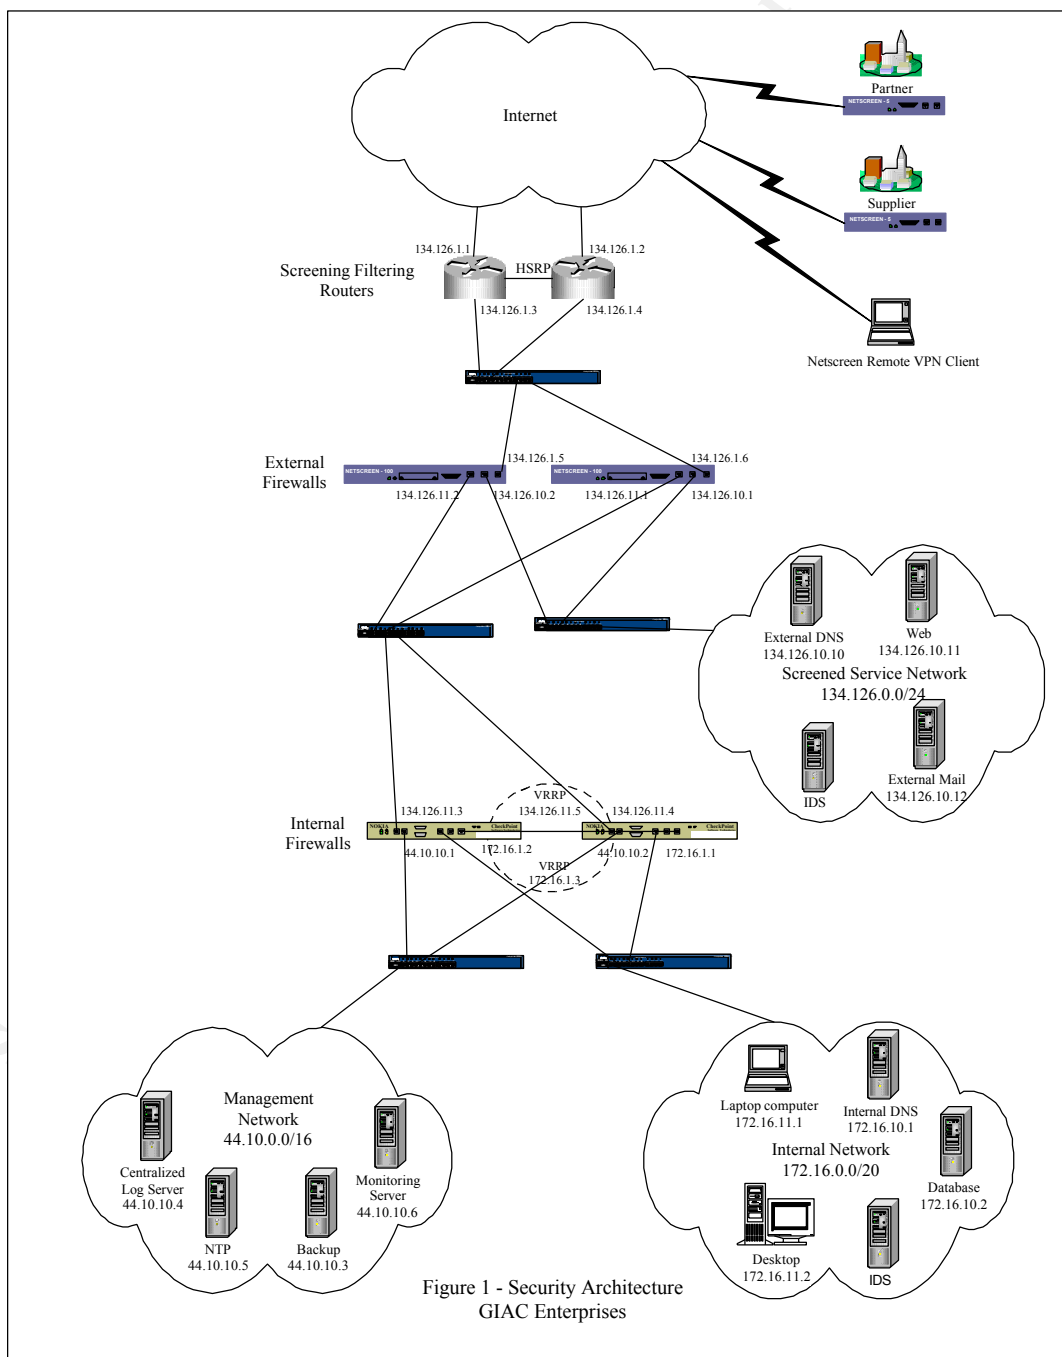
© SANS Institute 2000 - 2002, Author retains full rights.

4.0 ASSIGNMENT 4 – DESIGN UNDER FIRE

I have selected the design presented by Angela Orebaugh for this section. Her practical can be found at http://www.sans.org/y2k/practical/Angela_Orebaugh_GCFW.zip.

4.1 ARCHITECTURE UNDER FIRE

This practical was chosen because it very similarly tracks with several principles that I support. Her structure utilizes multiple vendor firewall types and out of band management. Her design is also a recipient of honors for its design so it is obviously a sound architecture and was explained very well in her paper. The purpose, however, of this part of the practical is to demonstrate that “perimeter systems are not magic ‘silver bullets’ immune to all attacks.” Her design, while very difficult to defeat, may not be immune. There are always areas that can be improved (as the attacks against mine in the future will also demonstrate) and an initially strong design may need adjustment over time.



4.2 ATTACK AGAINST THE FIREWALL ITSELF

Angela has chosen to use two firewall types to protect her architecture: Checkpoint and Netscreen.

The Netscreen firewall is a relatively newer entrant into the firewall market (compared to established names like Checkpoint, Cisco PIX, Raptor, etc.) and has few documented vulnerabilities (only 2, in fact). One is a firewall denial of service vulnerability and the other is a policy bypass vulnerability.

A vulnerability that allows a malicious attacker to crash the firewall due to the firewall's failure to handle exceptional conditions affected all Netscreen ScreenOS versions published January 8, 2001 on bugtraq (id 2176). This attack utilizes the same kind of excessively long URL attack used against many web servers. This was able to crash the firewall by exploiting its use of a local web server daemon for management that was vulnerable to this attack. The published exploit on SecurityFocus explains:

```
Once the input URL is longer than 1220 bytes=A3=ACNetScreen firewall= will crash:
$echo -e "GET /perl -e 'print "A"x1220" HTTP/1.0\n\n"|nc= netscreen_firewall 80
```

This exploit was unsuccessfully attempted against Angela's firewall. Netscreen did release new versions of the ScreenOS that answered this vulnerability. Angela's design was not vulnerable to it in the first place because she apparently was not allowing management connections to the external interface of the firewall.

Another vulnerability (bugtraq id 2523) can potentially allow a bypass of the firewall policy. This is potentially very damaging but no exploits have yet been found to take advantage of it. SecurityFocus explains:

Versions of ScreenOS, the inbuild OS of two models in the NetScreen line (NetScreen-10 & -100) contain a flaw which may permit some packets, of a type which has been denied, to enter the DMZ. The vendor notes that this vulnerability does not affect traffic reaching the protected network, and that this issue only arises under certain circumstances, dictated by traffic patterns on the network.

As a result of this vulnerability, potentially malicious packets of a type which has been prohibited in the device's policy may, to a limited extent, reach the DMZ network.

Since Angela is using the DMZ network, a potential future exploit may affect her architecture. It demonstrates that a perimeter defense cannot remain static, it must always continue to be worked upon (a lot like Castillo de San Felipe del Morro in San Juan continued to be built up even after initially completed).

Her other firewall is a Checkpoint firewall. Checkpoint is a widely used firewall solution and has numerous vulnerabilities that have been found over the years. Most are eventually fixed through new versions or patches supplied by the vendor. Some can also be overcome by architectural design. Any attack against the internal firewall will need to come from the DMZ because her Netscreen firewall

prevents connections from the Internet to the internal firewall directly.

The section on compromising an internal system is later, but since the design lumps the web servers, mail servers, and DNS servers all on the same subnet, vulnerabilities on any of these gives a free launch point against the others without a firewall in between. Separation of the DNS and mail servers from the web servers would be preferable since the web servers are permitted to communicate with the database servers. There may be exploits against web servers directly, but why exacerbate the risk by lumping the web servers with commonly exploited services such as external DNS servers and mail servers?

Given that a vulnerability is exploited on one of the protected DMZ servers (either through the allowed service on a BIND or sendmail buffer overflow exploit or some other means), the Checkpoint firewall could then be attacked. There are several vulnerabilities that are designed to bypass the firewall, but few against the firewall itself. One possible attack requires use of two vulnerabilities, one to bypass the firewall and another to attack it from the inside.

A denial of service vulnerability (bugtraq id 2238) makes it possible to cause the firewall to crash by exploiting a bug in the licensing manager. This attack would require bypassing the firewall and compromising an internal system (possibly using the RDP header bypass published in bugtraq id 2952) and then sending multiple spoof-sourced packets out through the firewall. This attack forces the firewall to add the IPs to the protected list. Once the license is exceeded (Checkpoint micromanages the number of protected hosts in their baffling licensing structure) the firewall begins to send messages to the console and begins to eat up CPU resources. After a while, the firewall will eventually hit 100% utilization and require a reboot. The RDP bypass vulnerability made use of the implied rule that allows any traffic through the firewall using RDP header information.

The exploit for RDP bypass was unsuccessful due to the fact that Angela had disabled the implied rules.

Angela's design makes it very difficult to attack the firewalls directly from the external side. Most of the possible vulnerabilities are mitigated by her architecture and use of a management network.

Achilles Heal

There is a vulnerability to Angela's design, however. It is designed to protect from only one direction, the Internet. She has an internal network as part of the design that may seem relatively secure upon first inspection, but it is the key to her undoing. Like her grouping of vulnerabilities on her DMZ, Angela has grouped several items on the Internal Network that really need to be separated.

The Internal Network is, among other things, a network that she has put the company's employees on. These employees can be used as pawns in the attack/defend chess game to weaken the overall architectural defenses. (Users are usually a weak link for security and often conspire in their own downfall. The secret is to make sure this doesn't translate into total capitulation of other resources as well.) This network is allowed through both firewalls, unhindered, to the Internet. There are no egress filters for traffic outbound from the Internal Network. What Ms. Orebaugh misses by allowing this is that the rule that allows the users out to the Internet also gives them free access to the Management Network. This hole increases the likelihood of several back door attack scenarios against the firewalls by

using the weakness of the Internal Network. (I'll get to the database server later.)

Users can (and do) download Trojan programs such as Netbus, Loki, etc. An attacker can then use this access to launch attacks against the Management Network. Once on the Management Network, exploitation of the firewalls is possible.

Social engineering can be used to focus the attack against GIAC Enterprises. The attacker may be able to convince a user to download the Trojan, give out passwords, etc. by posing as an IT engineer over the phone, email, etc.

Physical breaches are another possibility and much easier than many think. Most companies are not extremely security conscious about physical security. Even if they do pay attention to physical security, it usually does not extend far enough. Locking down access to the management network and servers is probably done. But, access to the user networks is probably much less stringent. An attacker could either pose as some kind of maintenance worker, marketing visitor, someone lost, etc. and gain temporary or permanent access to the user network (a wireless repeater plugged into a free LAN port in some unused corner or office).

From the Internal Network, access to the management servers can be accomplished using tools like Nmap to find the management server by searching for TCP ports 256-259, using sniffers/tools such as L0phtCrack to capture any userIDs/passwords, and attempting connections to the management servers using the gathered intelligence. Once on the management server, the firewalls belong to the attacker.

These risks can be severely mitigated by taking the following steps. Perform egress filtering from the Internal Network, dividing the Internal Network and separating users by function so that access to the Management Network can be more tightly controlled, and ensure networks/user devices with access to Management Network are given heightened physical security.

4.3 DENIAL OF SERVICE ATTACK

One of the most potent tools for denial of service is distributed denial of service (DDoS) attacks. These are almost unstoppable from an individual level. I have chosen the tool Tribal Flood Network for Win2K (TFN2K) as the attack tool. Unwary users load this tool as a Trojan on their systems and execute it as part of another program that was 'tainted'. The program then runs silently in the background until the master activates it. The master periodically polls the Internet, looking for slaves. Once enough slaves are found, the master informs all the slaves to flood a specific target. This master can be, in turn, controlled remotely as well so that tracing the link back to the actual attacker is extremely difficult.

My attack will be augmented by the fact that I will have the drones initiate a distributed smurf attack against the GIAC Enterprises sight. These drones will be told to ping the broadcast addresses of sites that are not set to prevent themselves from being used as broadcast amplification points (colleges are a prime target for this) using GIAC Enterprises IP address as the source address. This will allow the attack from 50 compromised hosts to be multiplied in magnitude until the echo reply traffic directed at GIAC Enterprises is more than the available bandwidth on their border (100 Mbps).

There is not way to completely prevent a TFN2K DDoS attack. DDoS is a problem for the Internet community as a whole to solve. The best course of action is to take steps to ensure that none of your hosts become drones, to filter the illegitimate IP ranges suggested by SANS, to configure your router to prevent your site from being a broadcast amplification point.

Some effects can, however, be lessened. ICMP traffic can be filtered. Devices such as the Captus filter can be used to dynamically filter. IP addresses can be changed if an attack is long term. None of these, however, are perfect answers and it usually comes down to throwing excessive bandwidth at the problem. Having an OC-48 to your site could protect you from these flooding attacks filling up your pipe, but that is a might expensive option.

If everyone does their part as an Internet community and limits broadcast amplification, performs egress and anti-spoofing filtering, and takes steps to prevent becoming unwitting drones, the flooding attacks will become harder to perpetrate.

4.4 ATTACK PLAN TO COMPROMISE AN INTERNAL SYSTEM

Returning to the highly exploitable Internal Network, this hole also leaves a very important system exposed: the database server. Unexplainably, the design combines the users on the same network as the database server. This is the most important piece of the architecture and needs to be protected above all things. It holds GIAC Enterprises only product: intellectual capital. It holds potentially damaging information: financial and personal information of customers and possibly the company. It holds the key to GIAC Enterprises future. It is on the same network as regular users, is routable to the Internet for outbound traffic, and has no egress filters in place to limit outbound traffic.

The same methods used to attack the firewalls by attempting access to the Management Network could be used instead to gain access to the database server (why bother attacking the firewalls when the prize is right there?). Then once on the database server, initiate a direct connection outbound to the Internet unhindered by egress filters or any breaks in routing and download the spoils to some remote repository.

This design also ignores the statistic that an overwhelming number of attacks are from internal users.

Like the discussion of the attacks on the firewalls, the risks can be mitigated. Egress filtering from the Internal Network, dividing of the Internal Network and separating users from the database and controlling the access, and putting the database into the architecture in such a way to prevent it from initiating a connection to the Internet directly.

4.5 CONCLUSION

Angela's design appears very difficult to penetrate from the outside. It is. Her design makes wonderful use of layered architecture, split DNS, multiple firewall vendors, out of band management, and other best practices. Her armor is strong, it is virtually impenetrable; except for one small chink. The problem is,

that chink is over her heart. It can be fixed.

The lesson to myself is this. My design also has a chink in it somewhere. I am not omniscient. It is the duty of the security engineer to never be satisfied with their design, to continually revisit it from different angles, enlist others to give fresh points of view, and to find that chink before an attacker does.

© SANS Institute 2000 - 2002, Author retains full rights.

5.0 REFERENCES AND ACKNOWLEDGMENTS

References (web accessible):

SANS Institute. "Consensus Roadmap for Defeating Distributed Denial of Service Attacks" Version 1.10 - February 23, 2000

URL: http://www.sans.org/ddos_roadmap.htm

SANS Institute. "Help Defeat Denial of Service Attacks: Step-by-Step" Revision: 1.41 – 23 March 2000

URL: <http://www.sans.org/dosstep/index.htm>

Senie, D. "Best Current Practice" RFC 2644. August 1999

URL: <http://www.rfc-editor.org/rfc/rfc2644.txt>

Carnegie Mellon Software Engineering Institute, CERT® Coordination Center. "Deploying Firewalls"

URL: <http://www.cert.org/security-improvement/modules/m08.html>

SANS Institute. "How To Eliminate The Ten Most Critical Internet Security Threats" Version 1.32 January 18, 2001

URL: <http://www.sans.org/topten.htm>

Captus Networks White Paper. "Protecting the Network from Denial of Service Attacks:

The Captus Networks TRaP Technology™" August 2001 v6.1e

URL: www.captusnetworks.com

ARCUS Data Security

URL: <http://www.contingencyplanning.com/arcus/index.cfm>

Snort - The Open Source Network Intrusion Detection System

URL: www.snort.org

Nmap – Security Scanning Tool

URL: www.insecure.org

Phoneboy FAQ. "Which Ports Does FireWall-1 Use?" 3 June 2000

URL: <http://www.phoneboy.com/faq/0105.html>

Phoneboy FAQ. "How Can I Disable Everything in Rulebase Properties?" 25 February 2000

URL: <http://www.phoneboy.com/faq/0345.html>

Orebaugh, Angela SANS GIAC GCFW Practical 6 March 2001

URL: http://www.sans.org/y2k/practical/Angela_Orebaugh_GCFW.zip

NetScreen ScreenOS Firewall Policy Bypass Vulnerability, bugtraq id 2523

URL: <http://www.securityfocus.com/bid/2523>

NetScreen Firewall Denial of Service Vulnerability, bugtraq id 2176

URL: <http://www.securityfocus.com/bid/2176>

Check Point Firewall-1 4.1 Denial of Service Vulnerability, bugtraq id 2238

URL: <http://www.securityfocus.com/bid/2238>

Check Point Firewall-1 RDP Header Firewall Bypassing Vulnerability, bugtraq id 2952

URL: <http://www.securityfocus.com/bid/2952>

King, Steven et al. Collaborative Task-Force funded by the NCES of the U.S. Department of Education
“Safeguarding Your Technology” Chapter 5 – “Protecting Your System: Physical Security”

URL: <http://nces.ed.gov/pubs98/safetech/chapter5.html>

Higgins, Scott. “Physical Penetrations: The Art of Advanced Social Engineering” SANS Practical. February 22, 2001

URL: <http://www.sans.org/infosecFAQ/audit/penetrations.htm>

References (Books):

Stevens, W. Richard. “TCP/IP Illustrated, Volume 1” Addison Wesley Longman, Inc, 1994.

McClure, Scambray, and Kurtz. “Safeguarding the E-Business Network” Excerpts from “Hacking Exposed” Cisco Press, 2000.

Nokia Systems. “Nokia VPN Gateway Configuration Guide” Version 3.1 August 2001

Nokia Systems. “Nokia VPN Gateway Administration and User Guide” Version 3.1 August 2001

Cisco Systems. “PIX Firewall Advanced” Version 1.01 2000

Sackett, George C. “Cisco Router Handbook” McGraw-Hill, 1999

SANS Institute. “Security Essentials” Track 1 Courseware 2001

SANS Institute. “Firewalls 101: Perimeter Protection with Firewalls” Track 2 Courseware 2001

SANS Institute. “Firewalls 102: Advanced Perimeter Protection and Defense” Track 2 Courseware 2001

Checkpoint Software. Security Courseware for Checkpoint 2000 Edition. 2000

Appendix A – Nokia Voyager Configuration Summary for External (Primary) Firewall

Nokia Voyager: External-FW Summary Tue Sep 11 19:43:44 2001 GMT

IP interfaces:

Physical interface	Speed	Duplex	Logical interface	Active	UP	IP address
eth-s2p1	100 Mbit	Full	Duplex External_Net	On	Up	208.a.1.2/28
eth-s2p2	100 Mbit	Full	Duplex OOB_Mgmt_Net	On	Up	10.1.100.1/24
eth-s3p1	100 Mbit	Full	Duplex PartnerReseller_Net	On	Up	192.168.10.1/24
eth-s4p1	100 Mbit	Full	Duplex Customer_Net	On	Up	192.168.20.1/24
eth-s5p1	100 Mbit	Full	Duplex VNP_Net	On	Up	192.168.100.1/24
loop0	loop0c0	On	Up			

DVMRP Tunnels

No interfaces

VPN Tunnels

No interfaces

ARP

Global Settings:

Keep Time

Retry Limit

Permanent Entries:

IP Address Type MAC
none

Routing Configuration

BGP

none

OSPF

No interfaces

RIP

No interfaces

IGRP

No interfaces

IGMP

No interfaces

DVMRP

No interfaces

PIM

No interfaces

Static Routes

Network	Next hop type	Next Hops	Priority	Description
Default	Normal	208.a.1.1		

Route Aggregation None

Inbound Route Filter Summary

Route Redistribution Summary

Routing Options

Equal-cost path : Maximum equal-cost path splitting: 8

Next hop selection algorithm: Source/Destination hash

Routing rank: OSPF routes: (default 10)

IGRP routes: (default 80)

RIP routes: (default 100)

OSPF ASE routes: (default 150)

BGP routes: (default 170)

Traffic Management

Access List Configuration

Default_drop Interfaces(Input):

Interfaces(Output):

Existing rules:	ID	Action	Src IP/mask	Dest IP/mask	Src Port	Dest Port	Protocol
	1	drop	0.0.0.0/0	0.0.0.0/0	0-65535	0-65535	any

Traffic Condition Configuration

No aggregation classes

Queue Class Configuration

Router Services

BOOTP Relay

No interfaces

IP Broadcast Helper/ Router Discovery

No interfaces

VRRP

No interfaces

NTP Not enabled

System Configuration

DNS

Not configured

Mail Relay

Mail Server: none

Remote User: none

Local Time Setup

Time zone: Etc/Greenwich

Host Address Assignment

External-FW On 10.1.100.1
localhost On 127.0.0.1

Change Hostname

Current hostname: External-FW

Manage Configuration Sets

Current configuration database: initial
Databases currently available: initial Tue Sep 11 19:42:30 2001

Configuration Backup and Restore

Backup files available:

Manage IPSO Images

Current selected image: IPSO-3.3-FCS8-05.14.2001-222300-651
Available IPSO images: IPSO-3.3-FCS8-05.14.2001-222300-651

Manage Installed Packages

Installed Packages: Check Point FireWall-1 (Strong) Version 4.1 SP-3
(Thu Feb 8 22:06:03 IST 2001 Build 41821) On
F-Secure SSH client and server, version 1.3.6.2 On

Security and Access Configuration

Passwords

admin S/Key Disabled
monitor S/Key Disabled

User Names

Username	Realname	Type
admin	Admin	Admin
monitor	Monitor	Monitor
root	Root	Admin

Network Access and Services

Access: Allow FTP access: disabled
Allow telnet access: disabled
Allow admin network login: enabled
Allow com2 login: disabled
Allow com3 login: disabled
Services: 'echo' service: disabled
'discard' service: disabled
'chargen' service: disabled
'daytime' service: disabled
'time' service: disabled

Voyager Web Access

Access: Allow web access: disabled
SSL security: enabled (168 bits minimum)
TCP Ports: Voyager (HTTP): 80
SSL Voyager (HTTPS): 443

SNMP Configuration

Read Community
n0pw4u2u\$e

Write Community
Not enabled

ColdStart traps: enabled
LinkUp/linkDown traps: enabled
Authentication traps: enabled
Frame Relay DLCI Status Change traps: disabled
VRRP Trap New Master traps: disabled
VRRP Trap Authentication Failure traps: disabled
System Trap Configuration Change traps: disabled
System Trap Configuration File Change traps: disabled
System Trap Configuration Save Change traps: disabled
System Trap Low Disk Space traps: disabled
System Trap No Disk Space traps: disabled
System location: GIAC Enterprises, Denver, CO
System contact: scot.hartman@GIACE.com

Monitoring Configuration

Collection Events Collection Switch Collection interval(Seconds)
Selected Report Type Selected Aggregation Class/Interface
Rate Shaping: On 60 Daily None
Interface Throughput: On 60 Daily None
Interface Linkstate: On 60 Daily None

Licenses

BGP: On
DVMRP with RIP: On
DVMRP with OSPF: On
DVMRP only: On
IGRP: On