Global Information Assurance Certification Paper

GIAC Practical Assignment for Capitol SANS

1. Security Architecture

Requirements

GIAC Enterprises is a small but growing business engaged in online Internet sales of information. The information sold, as well as customer data, is stored on an internal database server with strict requirements for availability, integrity and confidentiality. Public access to this data is mediated through an external web server which presents strict requirements for availability and integrity of the server and its contents, and for confidentiality of any data input by customers. The entire collection of GIAC Enterprises servers are dependent on the availability and integrity of information provided by infrastructure network services, in particular the Domain Name System (DNS) server(s). The GIAC enterprises security architecture is designed to enable authorized access to required information, while mitigating risks of unauthorized access, loss or modification of data, and denial of availability of critical services.

Architecture

GIAC Enterprises' security architecture, as depicted in Fig. 1, consists of a series of access control and intrusion detection devices arranged to provide defense in depth against external and internal threats. These devices include a perimeter router, an Internet firewall with Virtual Private Network (VPN) capability, and an internal firewall protecting the private server network, configured to create a layered defense. Access to the entire GIAC network is mediated by a filtering router, which is configured to reduce the risk of success of certain types of attack including such as IP address spoofing. In addition, all servers are protected by at least one firewall. To reduce the risk of failures due to misconfiguration and/or latent vulnerabilities in particular firewall products, servers containing the most sensitive information are protected by two firewalls, from two different vendors. To reduce the risk of undetected policy violations, whether from external or internal sources, network intrusion detection systems are installed on each layer.

The GIAC network is subdivided by the access control devices into multiple logical networks. These include a screened perimeter network, a private server network, and a corporate Intranet. Of these subnets, only the perimeter network uses public/routable IP addresses. The other networks use RFC 1918 ("Address Allocation for Private Internets") addressing.

Access by customers is to an SSL-enabled public web server located on the perimeter network. Access by suppliers and partners is to a private web server accessible only via VPN. VPN encryption between suppliers and partners is accomplished via one of two alternative encryption services, depending on supplier/partner capabilities:

-- Checkpoint SecureRemote VPN client to GIAC firewall or

-- IPSEC supplier/partner firewall to GIAC firewall.

The border router is a Cisco 3660 router running IOS 12.1. The Internet firewall is a Sun Ultra 60 running Checkpoint Firewall-1 4.1, also known as Checkpoint 2000, on Solaris 2.6. The Internal firewall is a Sun Ultra 10 running (Axent, now Symantec) Raptor 6.5 on Solaris 2.6. Intrusion detection systems are 300 MHz Pentium III Microsoft Windows NT 4.0 workstations running ISS Real Secure 5.0. Host based intrusion detection is also deployed on critical servers including all servers on the perimeter network, all servers on the private server network, all NT domain controllers and all DNS servers.

The architecture does not include any dialup access into the network. Where dialup access is required it can be accomplished using an Internet Service Provider, with the authorized client being provided VPN software.

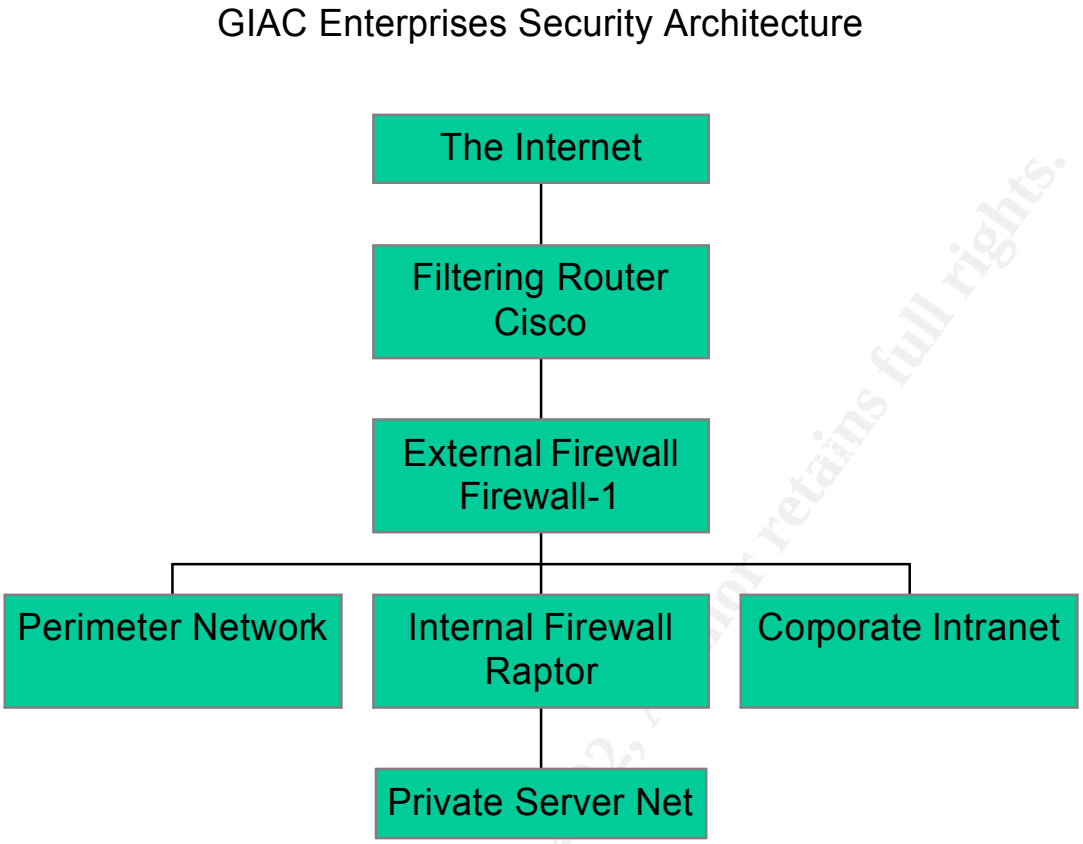

Fig. 1 GIAC Enterprises Security Architecture

Addressing

For purposes of the discussion which follows, assume GIAC enterprises has been assigned the class C IP address range 199.99.99.0-199.99.99.255. Further assume GIAC enterprise's network access provider has assigned an address of 199.88.88.88 as the external address of the GIAC Enterprises border router. These addresses are entirely arbitrary, selected for purposes of illustration only, and have no relation to actual real-world IP address assignments.

Using a subnet mask of 255.255.255.192 the GIAC address space can be subdivided into 4 subnets of which only two are populated initially in this architecture:

GIAC-External 199.99.99.0-199.99.63
GIAC-Perimeter 199.99.99.64-199.99.99.127

GIAC's private networks are configured as follows, using a subnet mask of 255.255.255.0:

GIAC-Private 10.99.99.0-10.99.99.255
GIAC-Intranet 10.99.100.0-10.99.100.255

Key systems are assigned the following addresses:

GIAC-BorderRouter: 199.88.88.88, 199.99.99.10
GIAC-ExtFirewall: 199.99.99.11, 199.99.99.75, 10.99.99.139 and 10.99.100.203

GIAC-NetIDS1: 199.99.99.1
GIAC-NetIDS2: 199.99.99.66
GIAC-NetIDS3: 10.99.99.130
GIAC-NetIDS4: 10.100.99.196
GIAC-ExtMail: 199.99.99.65
GIAC-ExtWWW: 199.99.99.67
GIAC-ExtDNS: 199.99.99.68
GIAC-ExtTime: 199.99.99.69
GIAC-Proxy: 199.99.99.70
GIAC-IntMail: 10.99.100.65
GIAC-IntWWW: 10.99.100.202
GIAC-IntDNS: 10.99.100.68
GIAC-IntSyslog: 10.99.100.69
GIAC-IntTime: 10.99.100.70
GIAC-IDSConsole: 10.99.99.71
GIAC-DB: 10.99.99.129
GIAC-App: 10.99.99.131

2. Security Policy

Policy

The basic security policy of GIAC Enterprises is to consider all information proprietary, and to deny access to it, unless access has been specifically authorized based on a business need to know the information.

Access is authorized in the following cases:

Unauthenticated users on the Internet, who represent potential or actual customers, are allowed to access the external web server, mail server, and domain name system server. Access will be allowed using only the TCP/IP protocols associated with web, mail transport, and domain name service query respectively.

Authenticated users, including suppliers, partners, and staff, are authorized to access specific internal web servers that present interfaces to private services, including corporate electronic mail, product submission, customer data, and the like. Authenticated access requires encrypted connections, using VPN client software (Checkpoint SecureRemote) or IPSEC firewall-to-firewall VPN, as well as presentation of individual authentication data, e.g. passwords.

Specific internal application and web servers, for example the public and private web servers, are authorized to access systems on the private server network using protocols including Oracle SQLNet.

IT support staff are authorized to connect to systems on the perimeter and private server networks using encrypted virtual terminal and file transfer services enabled through the use of the product SSH.

No users other than IT support staff and those with workstations on the finance and human resources subnet are authorized to access systems on those subnets.

Regardless of protocols used or other characteristics:

- No source-routed packets are permitted.

- No traffic crossing the external perimeter may use private addresses (IP address ranges 10.0.0.0-10.255.255.255, 172.16.0.0-172.31.255.255, 192.168.0.0-192.168.255.255, and the address 127.0.0.1).

- All traffic originating internally must have an internal address, and no traffic originating externally may have an internal address.

All network activity is subject to monitoring, banners/notices to this effect will be prominently displayed, and firewalls will allow network security access to the monitoring platforms.

No modems may be connected to the network without specific authorization. Violation is grounds for dismissal.

All systems connected to the network must be placed under configuration management, and upgraded as required to respond to vulnerabilities that are identified by vendors or reputable third parties.

VPN access whether individual (SecureRemote) or organizational (firewall-to-firewall) requires specific authorization and execution of a written agreement that requires the other party to maintain certain standards. Standards for individuals relate to maintenance of currency of antivirus and other protective technology, while those for organizations establish minimum standards with respect to their security policy and practices.

Implementation - Router

Router access control lists are used to enforce specific limited portions of this policy, while the firewalls enforce the remainder.

The router access control list entries required are as follows. Logging should be enabled to a local syslog server on GIAC-Intranet:

```
interface serial 0
ip address 199.88.88.88 255.255.255.0
no ip directed broadcast
ip access-group 11 in
access-list 11 deny 10.0.0.0 0.255.255.255.255 any log
access list 11 deny 172.16.0.0 0.15.255.255 any log
access-list 11 deny 192.168.0.0 0.0.255.255 any log
access-list 11 deny 199.99.99.0 255.255.255.0 any log
access-list 11 deny 224.0.0.0 31.255.255.255 any log
access-list 11 deny host 0.0.0.0 log
access-list 11 deny host 127.0.0.1 log
access-list 11 permit any

interface ethernet0
ip address 199.99.99.10 255.255.255.0
no ip directed broadcast
ip access-group 12 in
access-list 12 permit 199.99.99.0 255.255.255.0
access-list 12 deny any log
```

Implementation – Checkpoint

The firewall rules required to implement this policy on the Checkpoint FW-1 are as follows:

```
Allow Any to GIAC-ExtWWW 80/tcp Log
Allow Any to GIAC-ExtWWW 443/tcp Log
Allow Any to GIAC-ExtDNS 53/udp Log
Allow Any to GIAC-ExtMail 25/tcp Log
Allow GIAC-ExtMail to GIAC-IntMail 25/tcp Log
Allow GIAC-Intranet to GIAC-Proxy Any 80/tcp Log
Allow GIAC-Proxy to Any 80/tcp Log
Allow GIAC-IntDNS to Any 53/udp Log
Allow GIAC-IntTime to GIAC-ExtTime 123/tcp Log
```

Allow GIAC-Intranet to GIAC-Proxy Any 443/tcp Log
Allow GIAC-Proxy to Any 443/tcp Log
Allow GIAC-VPN-Encrypted to GIAC-IntWWW 80/tcp Log
Allow GIAC-VPN-Encrypted to GIAC-IntWWW 443/tcp Log
Allow GIAC-ExtWWW to GIAC-DB 1521/tcp Log
Allow GIAC-IntWWW to GIAC-DB 1521/tcp Log
Allow GIAC-Networks to GIAC-IntSyslog 514/udp Log
Allow GIAC-ExtTime to Any 123/tcp Log
Allow GIAC-IDSConsole to Any 2998/tcp Log
Allow GIAC-Intranet to GIAC-IDSConsole 901/tcp Log
Allow GIAC-Intranet to GIAC-Networks 22/tcp Log
Deny Any to Any any Log

The rules are ordered according to expected frequency of occurrence of matching traffic, in order to maximize performance.

The firewall's configuration/default settings also must be checked to ensure that settings external to the policy's rulebase itself do not enable traffic to bypass the rules outlined. These settings are contained in the firewall Policy "Access Lists" tab.

The policy rules are implemented on a Checkpoint Firewall by creating (or selecting predefined) database objects and services to represent the networks, systems, protocols, etc. covered by the policy, and by creating and ordering policy rules with respect to the objects.

An example object definition is depicted in Fig. 2. Example policy rules are depicted in Fig. 3.

VPN

VPN implementation using SecureRemote is on a per user basis and accounts ("user" objects) must be created in the database for any authorized clients. For each user an encryption mechanism must be selected, IKE or FWZ, Checkpoint's original mechanism.

Specifics of VPN implementation firewall to firewall vary depending on the details of the firewall used by the partner/supplier organization. However in general VPN requires definition of the encryption domain and encryption scheme(s) to be used, e.g. IKE. IKE configuration requires selection of encryption algorithm(s) e.g. DES and 3DES, hash algorithm e.g. SHA1 and/or MD5, and authentication method e.g. shared secret or alternative PKI certificates. Other IKE parameters are set in the firewall Properties Encryption tab, including time intervals for renegotiating IKE and IPSEC security associations.

Implementation – Raptor

The rules required for the policy on the Raptor firewall protecting the server network are:

Allow GIAC-ExtWWW to GIAC-DB 1521/tcp Log
Allow GIAC-IntWWW to GIAC-DB 1521/tcp Log
Allow GIAC-Networks to GIAC-IntSyslog 514/udp Log
Allow GIAC-IDSConsole to Any 2998/tcp Log
Allow GIAC-Intranet to GIAC-Networks 22/tcp Log
Deny Any to Any any Log

For brevity, screen prints of the implementation using the Raptor graphical user interface are omitted.

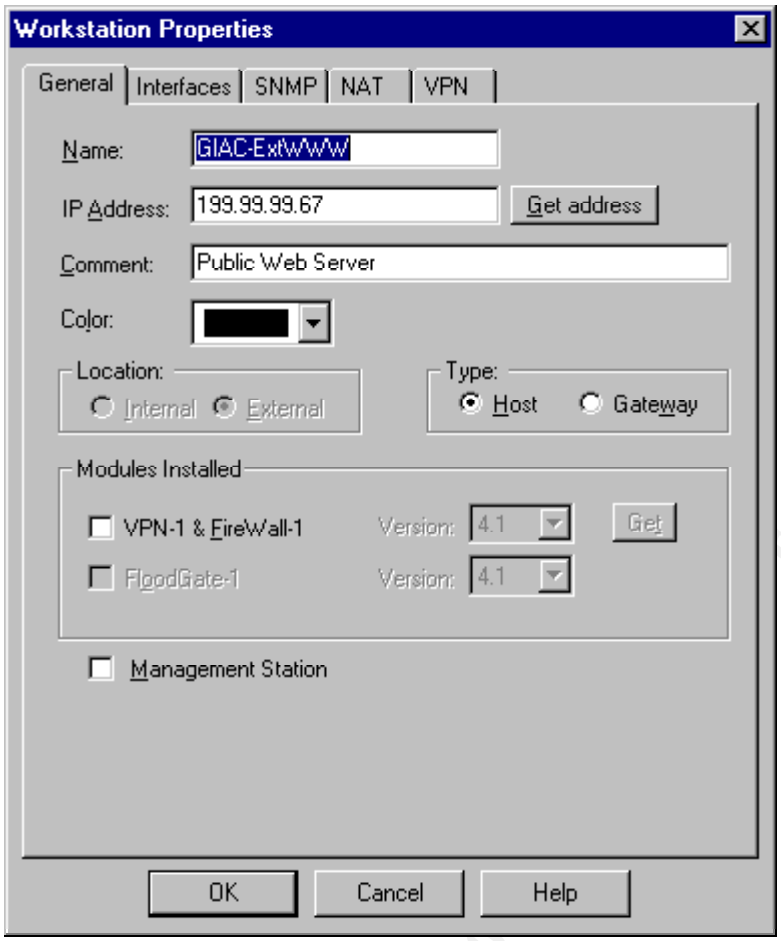

Fig. 2 GIAC Enterprises Public Web Server defined in FW-1

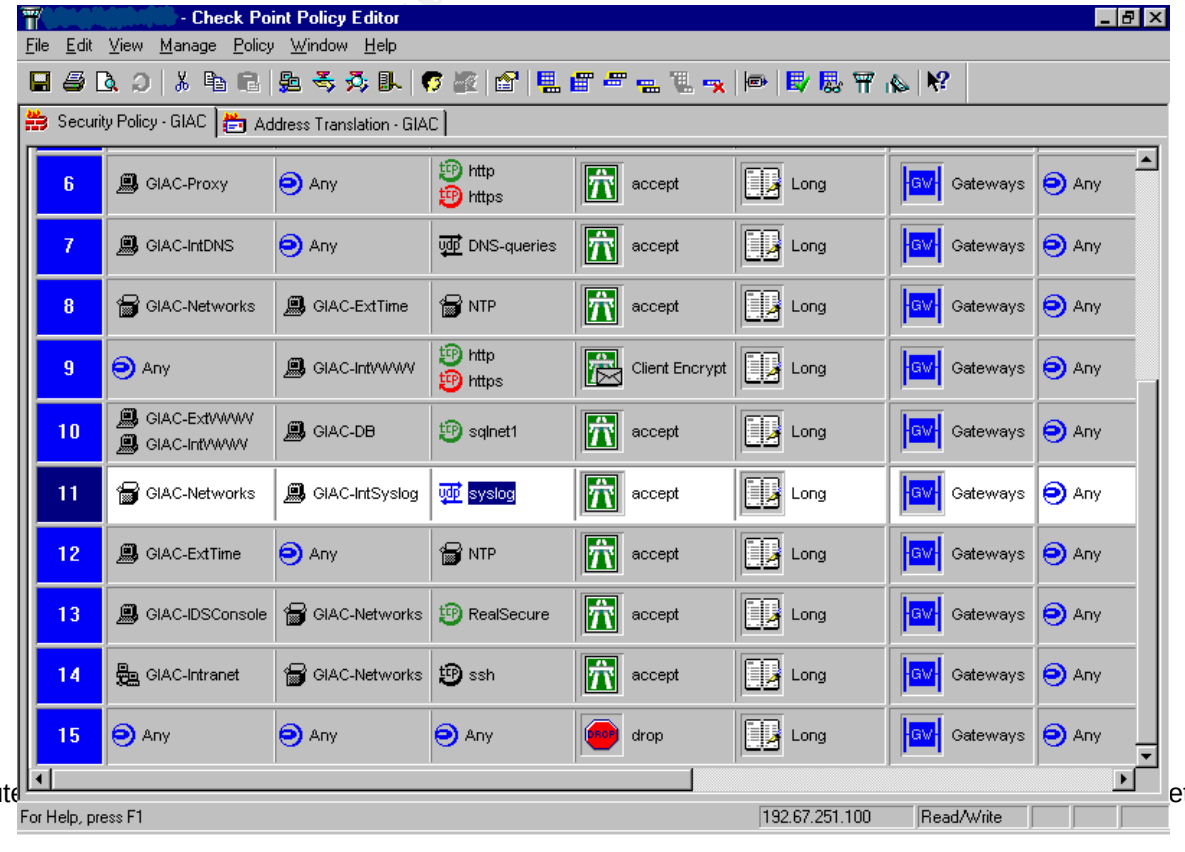

Fig. 3 Checkpoint Policy Editor Illustrating GIAC Rule-set

Risks

The security architecture, policy and implementation reduce but do not eliminate risks associated with network connectivity. Residual risks include risks of vulnerabilities in the commercial product components and/or their configuration, and risks associated with protocols allowed through the firewall.

Risks of vulnerabilities in components and/or their configuration can be mitigated by:

(a) assuring that staff responsible for the components are adequately trained in their use, (b) subscribing to services providing notifications of vulnerabilities and promptly responding to such notices, (c) budgeting for, scheduling and performing maintenance and periodic upgrades, (d) managing configuration changes, and (e) conducting periodic vulnerability testing.

In subsequent sections such practices will be referred to as "vulnerability management".

Residual risks associated with each protocol allowed through the firewall include the following:

HTTP: Risk is that the web content could be malicious. The risk is bi-directional. For example there is a risk that content from a client to a GIAC web server could cause a buffer overflow allowing unauthorized access, loss of integrity and/or confidentiality, and/or denial of service. And there is also a risk that content from an external server to a GIAC web client (browser) could cause unauthorized code to execute on the client, potentially turning it into an unwitting backdoor into the corporate network. Risks to web servers can be mitigated by vulnerability management and maintaining the currency of the web server product. Risks to clients can be mitigated by disabling or only partially enabling mobile code capabilities, and by maintaining the currency of anti-virus products.

HTTPS: Same as HTTP, except that in addition content is encrypted end-to-end, preventing inspection of content by IDS and/or proxy. Risks can only be mitigated by vulnerability management and close audit and review conducted on the web server.

DNS Query: Risk is that the DNS query or results could be malicious, and cause a buffer overflow or other event allowing unauthorized access, loss of integrity and/or confidentiality and/or denial of service. The primary risk mitigation locus should be the DNS server, where focus should be on vulnerability management and audit. (Also the DNS service should be replicated and hardened, however that is outside the scope of this document).

SMTP: Risk is that messages could have malicious content, and in particular malicious attachments, which could cause unauthorized activity on a client workstation, potentially turning it into an unwitting backdoor into the corporate network. This risk can be mitigated by subjecting the message contents to screening for malicious contents, with screening being done at multiple layers using multiple antivirus and related products. (Note however encrypted messages cannot be screened until the destination – usually a user workstation – is reached).

NTP: Risk is that the time service could be subverted, causing error or denial of service with respect to job/task scheduling and/or causing audit trails to be less/unusable and/or causing errors in authentication. This risk can be mitigated by using multiple time sources.

SQLNet: Risk is that unauthorized or malicious submission could occur causing loss of integrity and/or confidentiality of information or denial of service (as in the case of a submission causing excessive resource consumption).

Syslog: Risk is that unauthorized users could flood the syslog server with bogus logs, causing denial of service. This is mitigated by restrictions on sources authorized to "connect" via syslog, and by router filter

rules that prevent spoofing of source addresses.

RealSecure: Risk is that unauthorized users could identify and subvert IDS. This is mitigated by (a) vulnerability management with respect to the IDS and (b) use of strong encryption. Also in some a network architecture may be used that allows network IDS to use promiscuous-mode network interfaces to which no addresses have been assigned, and to communicate via a second, dedicated network.

SSH: Risks are that SSH authentication and/or encryption can be subverted, or that passwords may be guessed. These are mitigated by vulnerability management, and by user training and awareness.

VPN: VPN is mediated by shared secrets in most cases and there is some risk that such authentication keys could be guessed, if poorly chosen. However the major risk is that the VPN becomes an encrypted tunnel into the interior of the network, meaning that if the encrypted client or partner is compromised intrusions will only be visible after the decryption takes place. This is mitigated by placing IDS on server systems as well as on the network, by placing IDS inside the firewall, and by user training and awareness.

3. Audit

Planning

Auditing of the GIAC Enterprises network security implementation is accomplished through periodic online testing of the perimeter and interior access control devices using a commercial scanning product. The primary product used is the ISS Internet Scanner 6.1, updated to ensure that the most current attack patterns are applied.

Testing should occur at an off hour, at a time when traffic is normally lightest, here assumed to be 0000-0500 EST5EDT Sunday and Monday. Cost would be dependent on frequency of changes requiring re-testing. Prior to testing, all network and host-based intrusion detection systems on the network are examined to ensure they are operational, and time stamps in all router, firewall and IDS log entries are examined to ensure that they are accurate.

The first test performed is a full TCP, UDP and ICMP “port scan” of the entire address ranges of the perimeter, private server and internal networks, from the outside. Ports scanned should include the entire range from 0 to 65535, both TCP and UDP. The scanner log is examined subsequent to the scan, and should indicate no traffic getting through other than on those ports authorized for use. Router and firewall logs are examined to identify that the proper device is rejecting the traffic and recording the rejection. IDS logs are examined and should show detects of the network activity associated with the scan.

The second test is an individual network assessment of each perimeter protection device, using scanner configurations targeted at the underlying platform (Cisco IOS or Solaris). The scanner log is examined after the scan, and should indicate no vulnerabilities are detected in the configuration of the device. IDS and device logs are examined and should show detects of the network activity associated with the scan.

The third test is a full TCP, UDP and ICMP scan of the address ranges of the perimeter and private server networks, from the inside. The scanner log is examined subsequent to the scan, and should indicate no traffic getting through other than on those ports/to those destinations authorized for use. Router and firewall logs are examined to identify that the proper device is rejecting the traffic and recording the rejection. IDS logs are examined and should show detects of the network activity associated with the scan.

These tests are to be repeated on a regular basis, and in particular subsequent to any upgrade of the router and/or firewall software, and subsequent to any significant reconfiguration of the device.

Execution

The ISS Scanner supports a feature called X-Press Update, which is capable of downloading updated tests for the latest vulnerabilities. This feature should be used prior to every scanner session.

Prior to conducting any scan the scanner policy must be defined. This is ordinarily accomplished by selecting a predefined policy from those available from the vendor and using it as the basis for a scan tailored to the organization's needs.

Fig. 4 depicts some of the predefined scanner policies available. The L4 Router & Switch policy would, for example, be used as the basis for a local policy that would be used to scan the GIAC router.

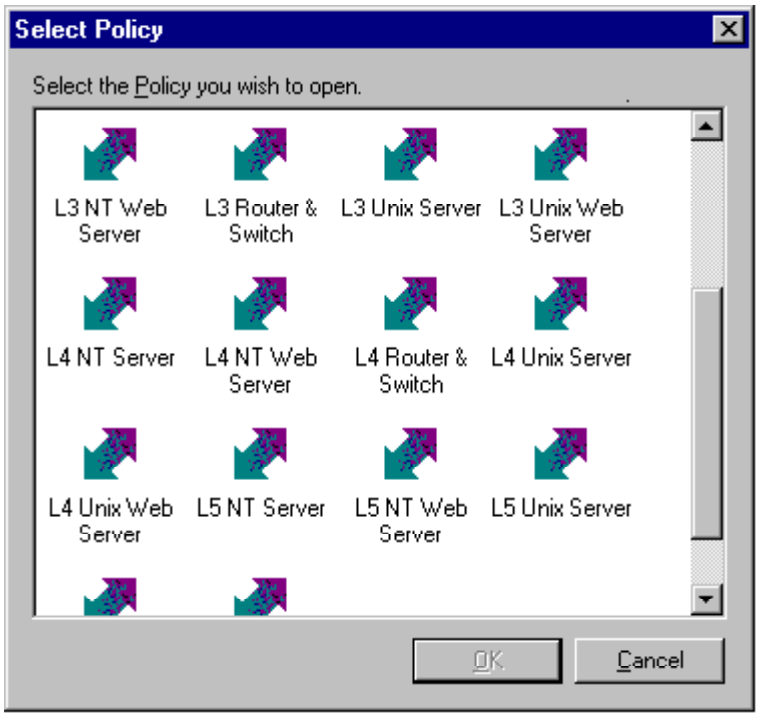

Fig. 4 ISS Scanner Predefined Policies

Editing of policies is accomplished by creating a new policy based on of the predefined ones, and editing it. Fig. 5 shows a screen of the Router and Switch policy open for editing.

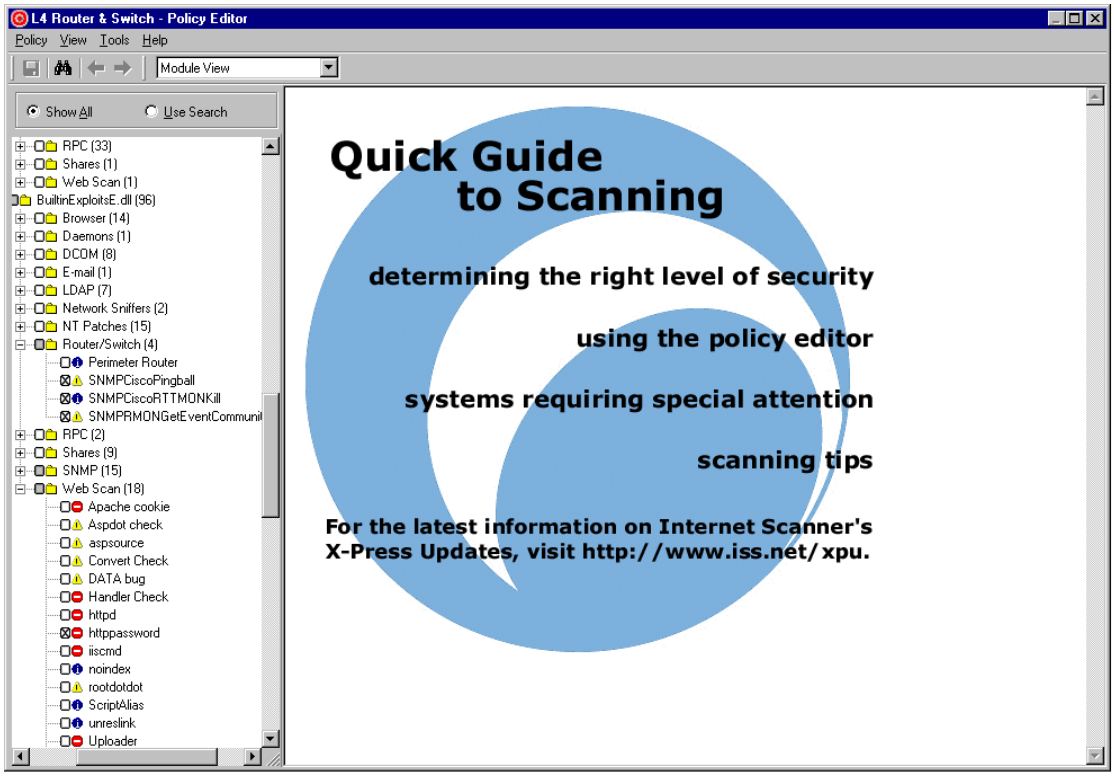

In the case of GIAC auditing, the only editing required is as follows:

a. Edit the Inventory policy to ensure that TCP and UDP scanning covers the full range of ports. Note: ISS port scanning tends to be slow, consequently GIAC may want to explore an alternative tool such as nmap for port scans to test firewall rule enforcement. ISS is however recommended for vulnerability scanning in any case.

b. Edit the various server policies to disable any brute force/password guessing attempts, as these tend to lock accounts. Instead of ISS Scanner, offline password cracking utilities should be used periodically to assess and enforce password quality.

To run the port scans, the auditor must adopt a variety of perspectives vis-a-vis the defensive layers. This is most easily accommodated if the scans are conducted using a portable computer. Specifically, using the terminology in Fig. 1:

a. The entire GIAC network address range should be port scanned from the Internet (with cooperation/permission of the ISP),

b. The entire GIAC network address range should be port scanned from between the Filtering Router and the Firewall-1, i.e. from an address in the range 199.99.99.0-63,

c. The private server network address range should be port scanned from an address on the Corporate Intranet,

d. The Corporate Intranet should be port scanned from an address on the Perimeter Network, and

e. The Private Server Network should be port scanned from an address on the Perimeter Network.

The vulnerability scans should be run from the network/subnetwork on which the device being scanned is located, i.e. avoiding any perimeter devices so that any vulnerabilities are not masked by the router or firewall.

After the vulnerability scans are run, reports are produced listing any services detected, vulnerabilities found or other results. The vulnerabilities are assesses as to whether they are High, Medium or Low risk. The reports include instructions covering how the vulnerabilities can be eliminated.

4. Design Under Fire

For purposes of discussion, the security architecture described at http://www.sans.org/y2k/practical/Heather_Bard_GCFW.doc was examined to determine what residual risks it might not have addressed. The architecture consists of a perimeter filtering Cisco router, and a Raptor firewall screening both a perimeter network and an internal network located behind a second filtering Cisco router.

Regardless of specific architecture, commercial products typically incorporate latent vulnerabilities requiring periodic patching or upgrade subsequent to purchase and installation. Thus known vulnerabilities associated with either the Raptor firewall or the Cisco routers or both might be present at any given time.

Any firewall policy permits certain traffic through, and as a result presents risks of vulnerabilities associated with the protocol(s) in question. In the case of the architecture and policy under review the protocols of most interest include HTTP and SMTP.

1. Attack against the firewall

Unpatched Axent Raptor 6.0 is susceptible to a malformed zero-length IP Options Denial of Service attack (see http://www.securityfocus.com, 1999-1021, Bugtraq ID 736). This is also described as Remote Denial of Service in Mitre's Common Vulnerabilities and Exposures, reference http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-1999-0905

Effect of the exploit would be to lock up the firewall causing a denial of service. Probability of success would be dependent on the currency of the patch level of the firewall.

2. Other denial of service

Among other potential vulnerabilities, Cisco IOS 12.0 and other versions can be crashed by malicious UDP packets to the syslog port, and can be rebooted by sending the ENVIRON option to the telnet service before it is ready to accept it. Also denials of service can be caused if maliciously crafted URLs are passed to the HTTP server. These vulnerabilities are identified as CVE-1999-0063, CVE-2000-0268, and CVE-2000-0360.

Even without the presence of vulnerabilities, router performance is also adversely affected by traffic volumes (and by patterns of filtered traffic that cause excessive router overhead). Flooding by TCP SYN, UDP or ICMP traffic would degrade performance either at the router or at the ISP, at thresholds dependent on the connection bandwidths and traffic volumes involved. One countermeasure would be requesting blocking by one or more upstream providers. Another would be having multiple network connections through multiple ISPs. The ultimate, slower, countermeasure is of course the criminal justice system.

3. Compromise of an internal system through the perimeter

The most serious sources of residual risk are the two protocols that nearly every firewall has to pass, SMTP mail and HTTP/HTTPS web traffic. Both of these protocol types involve externally generated content that is interpreted internally at a potentially improperly configured end-user workstation with results immediately apparent to a mostly untrained eye. An example of a difficult scenario:

a. An adversary establishes an unadvertised web server with active content, for example an Active X control that downloads and installs a remote control program that uses well formed HTML to send replies to the adversary and receive commands as HTTP replies.

b. The adversary selects a target by using Internet search engines to identify GIAC employees who have signed web server guest books, posted network news articles, or otherwise shown themselves to have "exploratory" tendencies.

c. The adversary then sends an electronic mail message to the target(s) inside the firewall which either causes automatic execution of the web browser or which entices the target user(s) to access the malicious site.

A variation on the above would be if the adversary would use Internet resources to identify business partners, and then impersonate the identity of a seemingly clueless staff member at the partner organization.

Jeffrey Roth
Capitol SANS Washington DC
December 2000 -> February 2001