Global Information Assurance Certification Paper

1.0 Security Architecture

1.1 Problem Statement

Assignment 1 - Security Architecture (25 Points)

Define a security architecture for GIAC Enterprises, a growing Internet startup that expects to earn $200 million per year in online sales of fortune cookie sayings, and which has just completed a merger/acquisition. Your architecture must specify filtering routers, firewalls, VPNs to partners, secure remote access, and internal firewalls. Be explicit about the brand and version of each perimeter defense component. Produce a diagram or set of diagrams with explanatory text that define how to use perimeter technologies to implement your security architecture.

You must consider and define access for:

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

1.2 Connectivity Requirements

While it is possible to develop an architectural design for a network by simply defining a set of components based on experience or typical installations, a robust design will generally be based on an understanding of logical and physical data flows.[1] This would normally be expressed in terms like connectivity, bandwidth, protocols, etc. For the purposes of this paper, the approach taken was to generate connectivity requirements by developing data flow diagrams. This required some, mostly implicit, process and organization assumptions to be made. Note that the flows and connectivity solutions selected and described below are not unique; many other variations are possible.

1.2.1 Direct product support

As shown in Figure 1, the first data flow considered was that of the primary product: fortune cookie sayings. For the sake of simplicity, "fortune cookie sayings" are hereafter generally called "verse." This notation would also be more useful when thinking about product diversification (.e.g., expansion into the very similar greeting card sayings business). Note one major assumption that has been made in generating this diagram—there will be two types of buyers. One category of buyer, the "online" buyer, is satisfied to buy non-exclusive verse using a credit

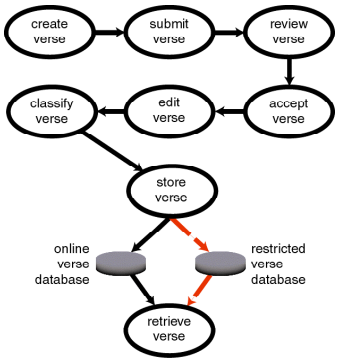

Figure 1. Product flow diagram.

[1] E.g., DeMarco, T., 1979, *Structured Analysis and System Specification*, [illegible], Englewood Cliffs, New Jersey.

card at published rates under terms of limited or restricted copyright privilege. The second category of buyer is one who is interested in special licensing or copyright agreements, exclusive use agreements, or possibly discount rates. The requirement to support exclusive use means that two verse databases must be supported with different sets of access and management controls.

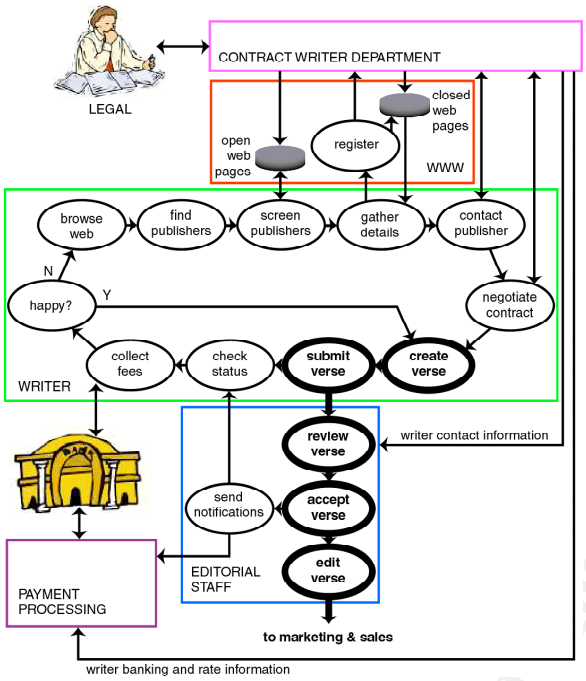

Figure 2. Product front-end detailed flow diagram.

The next step taken was to break down Figure 1 into several sub-flows and provide more detail. For the product front-end—the writing of verse—the results are shown in Figure 2. Some of the assumptions made during the creation of this and other diagrams include the following.

•

The problem statement provided in the assignment indicated that GIAC Enterprises is expected to earn $200 million per year. Assuming that a 3% profit margin applies, this sales figure indicates that something like 30 to 60 people are on the GIAC Enterprises staff. This is large enough to support an organization of multiple departments, as shown in the flow diagrams. Division or separation of duties and responsibilities among different departments is in keeping with basic computer security and accounting principles.[2]

- All writers work on a contract (out-of-house) basis, and are not "staff" per se.

- GIAC Enterprises is an "e-business" that conducts most or all of its transactions electronically.

- GIAC Enterprises is not large enough to have lawyers on its staff. Required legal support is sought from legal firms with appropriate expertise in e-business and publishing issues.

- Non-product data residing online is of two types: (1) open and freely available to all; and (2), open (non-sensitive) but restricted to registered users (as a means to gather contact information).

[2] E.g., Fisher, Royal P., 1984, *Information Systems Security*, Prentice-Hall, Inc., Englewood Cliffs, New Jersey, pp. 20-23.

- Sensitive, non-product data (e.g., supplier, customer, partner, or other contact information) is not kept in a manner that is accessible from the Internet.

- Use of central files servers and database systems would be possible. However, to provide the desired level of internal security, the cost of the security infrastructure (e.g., DCE/DFS plus DCE-enabled applications) was considered to be beyond that warranted for the current and expected near-term size of the company.

For the back-end of the process outlined in Figure 1, two diagrams were produced. One describes the 'online" buyer (a typical customer), while the other describes "contract" buyers—those with special requirements (e.g., a typical "partner").

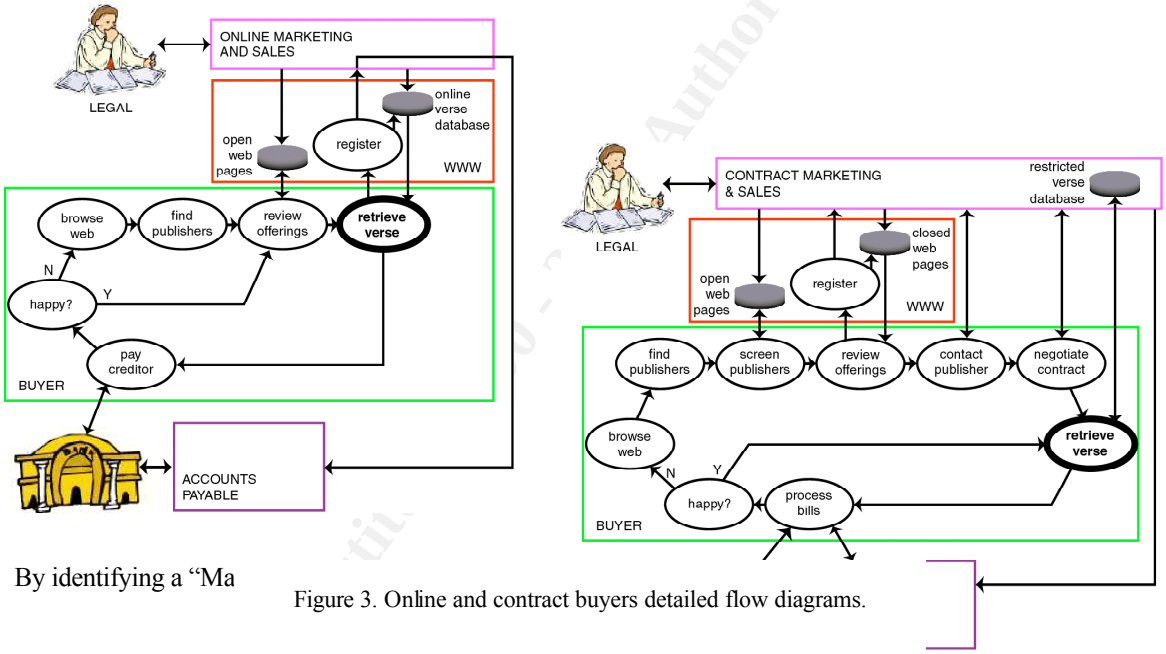

Figure 3. Online and contract buyers detailed flow diagrams.

By identifying a "Ma

3, that is also responsible for performing the remaining (middle) steps of the flow diagram of Figure 1, a connectivity requirements diagram can now be created, as below.

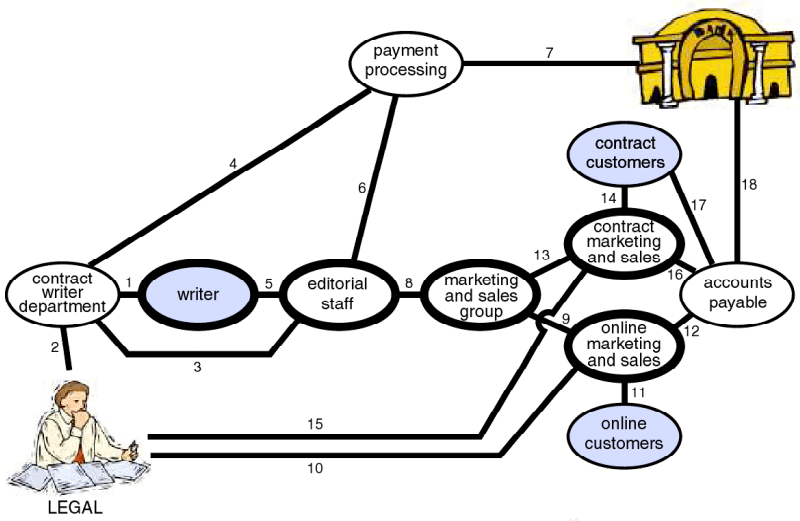

Figure 4. Basic connectivity diagram.

Comparison between the edges (connections) in Figure 4 with the diagrams of Figures 2 and 3, along with some implicit assumptions regarding applications and protocols, allows the following table to be generated (as noted above, other variations or configurations are possible).

Table 1. Connection Descriptions.

connection	**data exposures**				**protocols**			**routing**	
	Supplier data	Product	Customer data		email	Web	File transfer	Internet	Only on intranet[3]
			online	contract					
1	√				√	√		√	
2	√				√			√	
3	√				√		√		√

[3] The company intranet, as used here, may include VPN solutions over the Internet to bridge remote locations without the expense of private lines.

4	√				√		√		√
5	√	√			√			√	
6	√				√				√
7	√				√	√		√	
8		√			√				√
9		√			√		√		√
10					√			√	
11		√	√		√	√		√	
12			√		√				√
13		√			√		√		√
14		√		√	√	√	√	√	
15				√	√			√	
16				√	√				√
17				√	√			√	
18			√	√	√	√		√	

The first thing to note is that every connection but one is carrying some type of sensitive information (supplier or customer data , or product), and over one half of these are exposed to Internet-based threats. All of the connections are potentially exposed to intranet-based threats.

In recognition of the international nature of the business conducted by GIAC Enterprises (explicit international partners; implicit potential for international customers and suppliers), selection of safeguards to mitigate these exposures[4] must consider availability and distribution issues.

The obvious solution for exposure through email is to use PGP 6i.[5] At the same time, PGP provides the signature and non-repudiation support required for e-business. This version of PGP is available worldwide.[6] It can read and understand messages, keys and signatures created with PGP 2.0 and later. And it can generate messages, keys and signatures that can be read and understood by PGP 2.6.X and later, as long as RSA keys are used. PGP 6i is available for many platforms including UNIX, Windows and Macintosh computers. It can be used independently of the particular email client (or server) in use, although it has been integrated with many of the commercial email clients on the market today. It is cost effective (freeware versions are

[4] Note that we are here considering only the data exposures due to communication along the connection pathways. Security concerns due to data storage exposures are addressed separately.
[5] http://www.pgpi.org
[6] On 13 December 1999 the U.S. Government granted Network Associates, Inc. a full license to export PGP world-wide, ending a decades-old ban on the export of strong encryption products.

available—a useful consideration for "starving artists" (writers)). And, finally, all source code is freely available. It should be noted that such characteristics—cost effective, absence of design secrecy, and universal application—are considered to be safeguard selection principles.[7]

For data exposure on the web, the choice would be the HTTP protocol with SSL support (https). From a client perspective (suppliers and customers), many browsers are internationally available on a wide range of hosts that support this protocol, most even as freeware or packaged as part of some other (e.g., OS) offering. For the concerned client, the source for some browsers is even available. From a server standpoint, cost effectiveness (free, although commercial support is available) and open source code concerns are best met with Apache and OpenSSL or SSLeay.[8] Certificates are available for "Apache-SSL" from over a dozen providers, or they can be self created from tools that come with OpenSSL and SSLeay.

In terms of securing file transfers, the option of choice is to use the secure file copy (scp) command available as part of the secure shell (ssh) package.[9] A wide variety of platforms are supported, with both free and licensed versions, and source code for some options.

Other VPN solutions (e.g., appliance using IPsec) might be possible, but their easiest use is probably for internal company applications. Examples would include connectivity between the GIAC Enterprises main office and the merger/acquisition site, or between GIAC Enterprises employees on travel and their home office.

Other options are available for reducing the risk of information on the intranet in addition (**not** in lieu of) to the protocols selected above. This would include: (1) restricting the content of information transmitted to need-to-know (e.g., as implied by Figure 2, neither the editorial staff nor payment processing needs all, or even the same, information concerning a particular writer; they only get the minimum required to perform their job functions) as documented in appropriate policies and procedures; (2), providing connectivity between departments only as required (e.g., for product support as developed above) and as approved in appropriate policies; and (3), providing physical segregation of signals where reasonable (and not just virtual—that might be hacked).

An example of the last option—physical segregation (related to the computer security principle of division or separation id duties discussed earlier)—is afforded by the recent merger/acquisition by GIAC Enterprises and by the nature of the data flow (the two "halves" can be spilt at connection 8; see Table 1). By placing one-

[7] Parker, Donn B., 1981, *Computer Security Management*, Reston Publishing Company, Inc., Reston, Virginia, pp. 170-183.
[8] http://www.apache-ssl.org See also http://www.thawte.com/support/server/apachessl.html Another option is to use mod_ssl with Apache.
[9] See http://www.ssh.com http://www.openssh.com and http://www.freessh.org

half of the business at the merger/acquisition location (including use of legal and banking support at that location), the only data flow between the two is the forwarding of product from the Editorial Staff Department (ESD) to the Marketing and Sales Group (MSG). (Although it is reasonable at this point to assume that the MSG will transmit requirements and performance information back to the ESD, say via email; this would communicate the type of verse needed (e.g., selling well or projected market), and may be used, in turn, by the ESD in their selection of verse from that submitted by the writers, etc.) The ESD need have no knowledge of actual customers and the MSG need have no data concerning the writers.

1.2.2 Corporate indirect support

While it would be nice to think that connectivity requirements have been neatly "sewn up," the fact is that much more is needed to provide the necessary indirect support to the departments discussed above simply to make a business operate. Some categories of indirect support that might apply are: computer support; server administration; purchasing; security; travel; financial auditing; human resources; payroll; and upper management. For the purposes of this paper, these functions have been gathered into the groups described below (and note that each group will have a presence, call it group -2, at the acquisition/merger site).

Services Group (SG and SG-2)—responsible for telecommunications and networking operations; server administration support; network services (e.g., middleware services, like name resolution and time services, mail server(s), and an internal corporate information server); computer support; computer security operations; safety and physical security; facility management and operations. Connectivity to all departments by email (IN/OUT), http/https (IN; corporate server), and ssh/scp (IN; e.g., to retrieve software updates or to update corporate server). Will also require outgoing Internet connectivity (email/http/https/ftp) in order to acquire software patches and upgrades; may be able to restrict these connections to specific destinations.

Business Management Group (BMG and BMG-2)—responsible for payroll, procurement, corporate contracts, expense reimbursement, spend planning, logistics, management reporting, and oversight of the Payment Processing and Accounts Payable Departments. Connectivity to all departments by email (IN/OUT) and http/https (IN; front end to ORACLE or other commercial business management software). Will also require general outgoing Internet connectivity (email/http/https) in order to effect purchasing, banking, and legal support.

Human Resources Group (HRG and HRG-2)—responsible for staffing, compensation, benefits, employee relations. Connectivity to all departments by email (IN/OUT) and http/https (IN; front end to PeopleSoft or other commercial HR management software). Will also require general outgoing Internet connectivity (email/http/https) in order to contact, for example, legal and staffing services.

Small Staff (SS and SS-2)—upper management administrative and secretarial support, and auditing functions. Connectivity to all departments by email (IN/OUT). Will also have to provide outgoing Internet connectivity (email/http/https/ftp) because they will chaff under restriction?!

Travel—arrangements are made through a web-based service under contract to GIAC Enterprises. Corporate employees will require outgoing Internet connectivity (email/http/https) in order to access travel services.

It should be noted that just as the communications for the two business halves should be segregated (departments dealing with the writing end from those selling the verse produced by the writers), the corporate indirect network should be segregated as well. This notion will give rise to three principal networks within the business, each with possible subnets for each different organizational element.

1.2.3 Other connectivity requirements

Provide remote connectivity for employees from home or while on travel. Without any analysis of risk, it is assumed here that this connectivity will be restricted to corporate services (e.g., email, corporate server, human resources, and business management servers). This is to account for the reduced level of trust of the security of the remote machine being used for access (as opposed to developing tight oversight and control of these resources—which is probably not very practical or effective anyway).

1.3 Asset Protection Requirements

A complete security analysis will include an assessment of risks from threats such as natural disaster, loss of physical plant (e.g., power and air conditioning), etc. For the purposes of this paper, however, the focus will be on the protection of information from cyber threats. The assets are information that resides on various computers across the company. In order of importance, that data is:

1. Supplier and customer data (e.g., names, contact information, and banking information). Verse supplier data resides in various forms within the Contract Writer Department (CWD), Editorial Staff Department (ESD), and Payment Processing Department (PPD). Verse customer data resides within the Online Marketing and Sales Department (OMSD), Contract Marketing and Sales Department (CMSD), and Accounts Payable Department (APD). Other supplier data resides within the Business Management Group (BMG and BMG-2).

2. Product data (verse or fortune cookie sayings). Resides in the ESD, Marketing and Sales Group

(MSG), OMSD, and CMSD, including the MSG WWW servers.

3. Other business data (e.g., payroll and other banking data). Resides within the BMG.

4. Human resource data (e.g., employee records). Resides within the HRG.

5. Other computer resources with very limited or no "sensitive data" (categories 1 through 4 above), but containing purchased and in-house software (e.g., desktop and workstation computers, middleware servers, internal and WWW corporate information servers, and dialup or remote access servers). Servers are maintained by the Services group (SG). Other assets are distributed across the company. Certain non-sensitive but controlled corporate data published by the CWD (on the CWD WWW servers) and by the CMSD (on the MSG WWW servers) also fall into this category.

In general, then, safeguards ought to be applied in consideration of this ranking. More protection should be applied to the "1" end, and less to the "5" end. One approach to protecting the high-value assets (servers) would be to isolate them with their own application-gateway firewall. However, this leaves open the question of how to protect the data that is retrieved from the server by an authorized client. Since the work flows defined for GIAC Enterprises allowed implementation of a fairly well-defined and compartmentalized organization, and since it is assumed that someone assigned to work in a department of this organization has general access to all of the respective data in order to perform their job function, the approach taken here is to use a firewall to protect all of a sensitive department's assets. Policy and auditing will be used to help keep data aggregation risks on department clients at a generally acceptable level, with the firewall providing the needed assurance that the data is not being unnecessarily exposed to external (to the department) threats. (In addition, a host-based packet-filtering firewall product will be used on the sensitive-data server to provide added defense-in-depth.) Packet filtering can also be used on routers located "in front" of firewalls to help control traffic, and thus provide some additional protection.

This is not to say that the "low-value" (from an information point-of-view) assets should not be provided some level of protection. While internal threats may exist, once a connection is made to the Internet, very-large numbers of potential threat agents can now "beat on the door." Corporate assets are then, in general, protected from the Internet by restrictive firewalls. However, in order to provide WWW-based and remote login services, some assets (e.g., HTTP servers) generally have to be placed "outside" of any such restrictive firewalls; less restrictive packet-filtering firewalls or "border" routers can be used to help minimize the threat to these machines.

The end result of these considerations is a tiered architecture that provides defense-in-depth for information

assets. Each of these network devices will have one or more security features in play that must be compromised before a particular threat can gain access to some designated target. (Note this is a general observation on the architecture discussed above, since no actual devices, security features, or rule sets have been defined.)

1.4 Applicable Technologies

1.4.1 Security components and nomenclature

The third element needed for developing a network-security architecture is an understanding of the security technologies or functional components available for deployment. Simply calling for a firewall is not sufficient, because the noun *firewall* is used in many ways these days. At best it means nothing more than[10] "software used to deny access to a network from outside that network, whether the access in question is through a direct login or through the use of program files uploaded to that network." The question is one of functionality, as outlined below. Note that the discussion is focused on functionality and not hardware or software packaging (that may implement multiple functions as defined below).

packet filters—The original firewall. Decisions to forward (route) or drop packets between a client and server are based on data contained in the network and transport headers (e.g., IP addresses and transport protocol). The more advanced packet filters include additional features that manage fragmented packets (e.g., complete reassembly) before forwarding decisions are made. Example packet filtering protection tasks include:

- *Block unauthorized IP addresses.* This would include: (1) unexpected or unassigned addresses that might be exploited, such as local addresses and the private addresses defined in RFC 1918;[11] or (2), IP source addresses appearing on the wrong interface (spoofing). Known bad or undesirable addresses could also be added to this list.

- *Block unauthorized services.* If a data flow analysis has been conducted for a network, the protocols and services (ports) authorized to pass between subnets should be well understood; those should be the only ones configured to be passed across a router. To help protect a network against certain mapping and denial-of-service (DOS) attacks, the router should also be carefully configured as to how it handles control messages; e.g., source routed packets or direct broadcast requests should probably be dropped,

[10] http://coyote.csusm.edu/public/guests/history/netinfo/firewall.html
[11] http://www.rfc-editor.org/

and use of return ICMP error messages should be carefully controlled.

dynamic packet filters—Some advanced packet filters understand the concept of the virtual connection that exists between a client and server, and make routing decisions based on this understanding; potentially this could be done for any transport protocol. (Stateful inspection or dynamic packet filtering; e.g., incoming packets have to match an original outgoing connection request or they are dropped, which requires the filter to have a detailed understanding of the protocol, and possibly application, that is in use.)

proxy servers—Also known as a diode. Used to prevent direct links between a client and server, and thus can be used to control and bring certain types of security elements to a connection. Clients must first connect to the proxy, which then connects to the desired server and passes the data through (both directions). Circuit-level proxies operate at the transport layer of the OSI model, while application-level proxies (aka. application gateways) operate, as the name implies, at the application level. By their very nature, proxies implemented on dual-homed hosts "hide" internal addressing (much like a Network Address Translator or NAT). Proxies can implement authentication and authorization for connections. Application proxies can implement content screening. Each supported transport or application protocol requires its own proxy.

reverse proxy servers—Proxy servers are designed to allow and control specific types of outgoing connections (those that originate from behind a firewall and go out to a server in a higher threat-level network such as the Internet). With the advent of web servers and e-commerce, it is also necessary for most organizations to provide services to clients on the Internet—which means servicing connections that originate somewhere outside of a firewall. To avoid a deliberate "breach" in a network design, servers accessible by clients on the Internet are generally placed outside the primary firewall. At a minimum, servers fulfilling this role should be "hardened" in order to minimize vulnerabilities to Internet-based threats (servers so hardened with service-specific configuration requirements are generally known as bastion hosts). However, it is possible to further reduce application-specific vulnerabilities by using proxies for *incoming* connections—so called reverse proxies. The most commonly thought of reverse proxy application is a server that manages web requests (http/https) for and in front of a content server; a breach of the reverse proxy will most likely only provide the perpetrator information involved with a single transaction, as opposed to having access to the entire database. Another good example of this idea is found in Simple Mail Transport Protocol (SMTP) "wrapper" programs (although it is rarely called a reverse proxy since it can be used to handle both incoming and outgoing "connections"). For both web and email applications, the ability of a proxy to inspect

content also adds a great deal to the security of a system by allowing harmful commands to be filtered out (e.g., email virus) before the data ever reaches the real server.

traffic-security devices—While information authentication (e.g., signature or non-repudiation) and encryption was discussed in section 1.2.1 above (e.g., the use of PGP at the application layer), it is also possible to provide security services at the network[12] layer. Two solutions are considered here: the use of strong authentication (two-factor; e.g., SecurID or other token) to enable connectivity to restricted services for remote clients via a Network Access Server (NAS); and the use of Virtual Private Network (VPN) technologies (e.g., IPsec) for data authentication and encryption while transiting the Internet. It should be noted that NAS typically requires the support of one or more security servers (SS) in order to handle the processing of token challenge and response protocols.

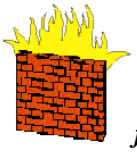

firewall—A network device employing one or more of the security components discussed above. Historically this included packet-filtering routers, and today it still includes computing devices that exclusively employ only packet filters (e.g., a dual-homed computer using *ipchains* under Linux). The name applies equally as well to machines running proxies (at least dual- or multi-homed systems). The best (or at least most flexible) of the firewall solutions will, however, provide both packet-filtering capabilities and various proxy applications as standard fare. Many firewall products also offer integrated VPN and NAT solutions as well.

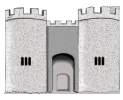

security gateway—A collection of security components, including firewalls and servers, that provides the desired connectivity between an internal network and the Internet in a secure manner.

1.4.2 Other network components and nomenclature

routers—Designed to provide connectivity between different subnets on the basis of data

[12] Security devices that operate at the data link and physical layers are also available, but were not considered for the purposes of this paper.

packet header information (OSI layer 3). However, over recent years, router device vendors have continually expanded the capability and flexibility of their devices. As alluded to above, the first extension was to add a security feature that made forwarding decisions based not only on the destination address (i.e., send to port X), but on the source IP address as well (i.e., is this packet source IP allowed to talk to the requested destination IP on port X; e.g., the Cisco standard access control list (ACL)). A related security feature was the addition of a logging capability. The inspection of source IP was soon followed by other straightforward extensions that analyzed the remaining fields in the IP header, and then expanded into the transport headers (e.g., the Cisco extended ACL). Now routers are available that build state tables (e.g., the Cisco reflexive access list) and check application layer information (e.g., Cisco Secure Integrated Software with Context Based Access Control), which essentially gives them the capabilities of many firewalls. And there are integrated switch and VPN solutions as well. The approach supported in this paper is generally to maintain separate devices for the different functions that must be performed (let a switch switch, router route, etc.), adding in only some basic security options rather than making full use of the advanced capabilities now available. However, the ability to perform basic (packet filtering) protection tasks enable a router to perform duty as a "front line" defense mechanism, or border router, by providing controlled connectivity between the Internet and all other company resources. This same ability will also allow a router to serve as an internal defense mechanism, or choke router, by providing controlled connectivity between different corporate assets. The use of multi-port routers can further strengthen a cyber security program by promoting traffic segregation. Note also that security should be applied to all traffic, regardless of the direction it is flowing across a router. The rules with most likely be different for the different directions, but it is unlikely that *no* rules would apply. (Also note that the routers themselves must be secured or hardened, but that is another topic.)

switches—Originally designed to provide connectivity between devices on the same subnet on the basis of data frame header information (OSI layer 2). However, with advances in technology, it is now common to find switches that operate at layer 3 as well. Switches can promote cyber security by supporting traffic segregation. Port security can also be used to help protect against address spoofing or the use of unauthorized equipment. For example, where tight controls are available on the installed equipment, the use of a Ethernet switch to provide connectivity will help prevent a threat agent from launching an attack from a compromised system on the local network using the IP address of a different computer (the switch, such as a Cisco 1900 series, would be configured to lock down the port based on an IP address violation). In a similar manner, a layer-2 switch can be used to lock out a port that sees an unauthorized MAC address broadcast, which might be appropriate for areas where tight access controls are not available.

access server (AS)—Designed to provide connectivity between a public switched network

(PSN) and a local area network (LAN). Think modem bank. As noted above, an AS can provide a strong authentication entry point into a network for remote, dialup users. The AS can also provide protocol screening (think ACL) much like a router.

1.5 Architecture Description

1.5.1 Internal network cyber security elements

The requirements and assumptions made in the previous sections on connectivity and asset protection can be collected together and transformed into an internal network design, as shown below in Figure 5. In this figure, "main office" refers to the primary GIAC Enterprises corporate location, while "remote office" refers to the offices associated with the merger/acquisition. The organization abbreviations used are as defined in sections 1.2 and 1.3 above. The principal protective feature in the internal network design is the use of segregated data (by department and group) that are protected by firewalls (packet filters and proxies as appropriate for the required connectivity; the security gateways, that also include the use of firewalls, are discussed in the next section). Data security is also supported by performing traffic segregation through the use of multi-port routers, and through the use of VPNs (as implied in Figure 5) to provide intra–departmental and –group connectivity between sites. And, although there are certainly ways around such precautions, it is envisioned that layer-2 switches with port locking will be used to provide machine connectivity to the network. As discussed in section 1.2.1, for the current and near-term size of GIAC Enterprises, a central security server for authentication and authorization services does not seem warranted (although this is not precluded by the selected architecture). Rather, password- and certificate-enabled accounts for restricted corporate services (i.e., those provided by the Business Management and Human Resources Groups) will be managed manually. The final element, that is not shown in the diagram, is the use of network intrusion detection (NID). Each principal security component (e.g., firewalls and gateways) should have a NID located behind it (the more secure side). These NIDs should be individually configured to specifically monitor for indications that the security component it is monitoring has failed or been compromised (e.g., rule failure). (While it is possible to train and employ analysts to look for unusual signatures in the packets transversing the network in "real" time, and while it is certainly true that NID logs can be vitally useful in reconstructing events post compromise, it is felt that the principal importance of a NID lies in its ability to instrument and verify the proper performance of network security components.)

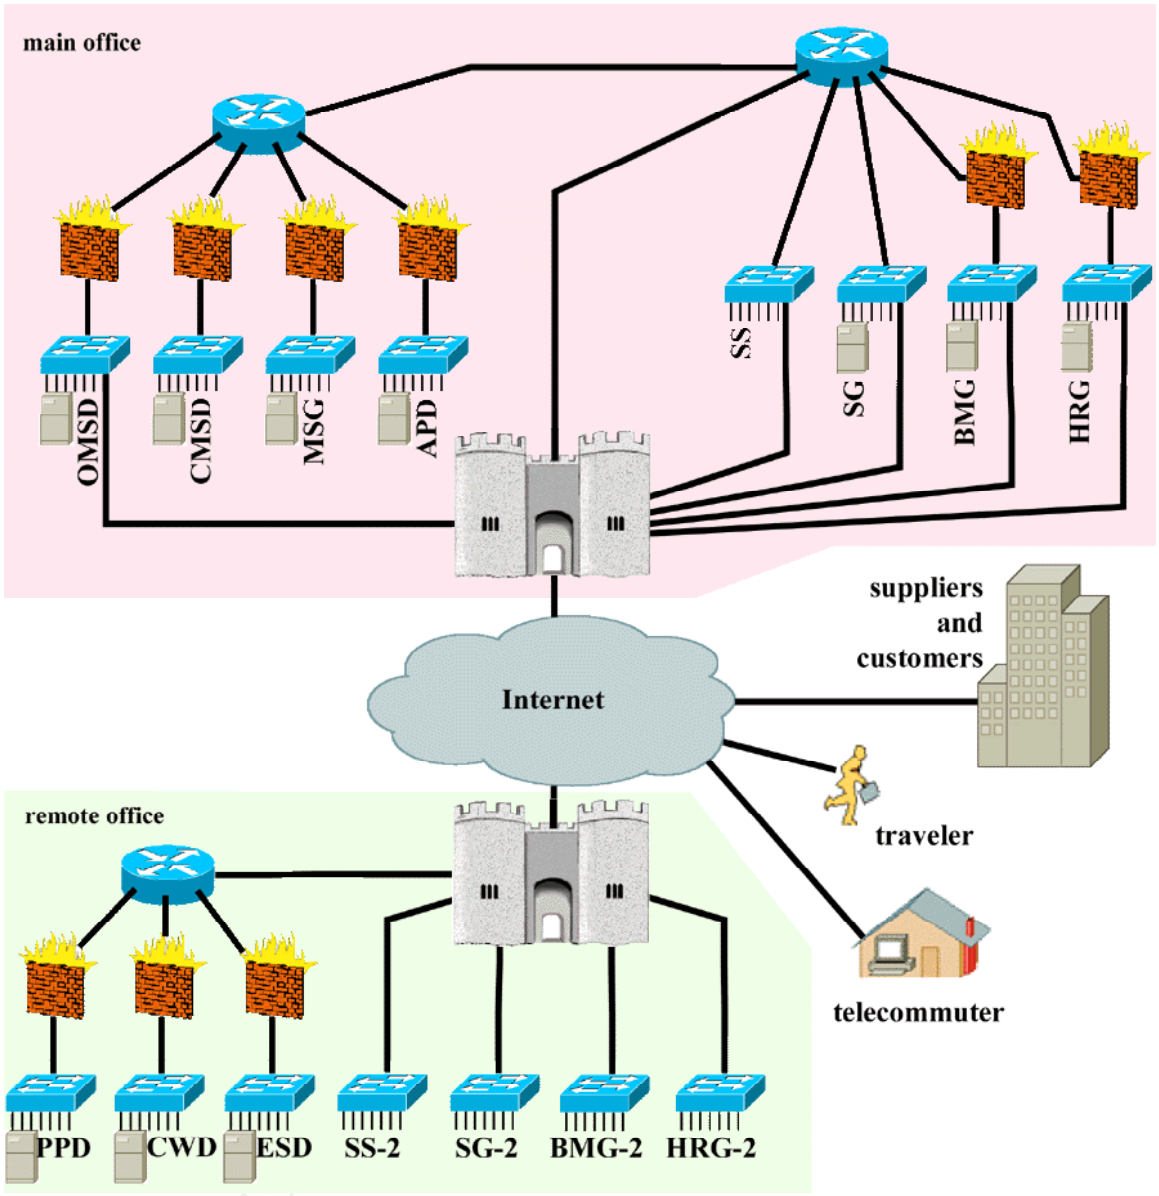

Figure 5. GIAC Enterprises network diagram

1.5.2 Network perimeter security elements

As for the design of the internal network, connectivity requirements established in the preceding sections can be used to lay out the design for the security gateways needed by GIAC Enterprises. The result, as shown below in Figure 6, is herein expressed in terms of the required security functionality (components); selection of the appropriate hardware and software elements to implement this design is discussed below. Note that while the discussion below is kept general, not all elements or data are present in both gateways (e.g., the OMSD WWW server and firewall, elements 27, 29 and 31, are only found in the main office gateway).

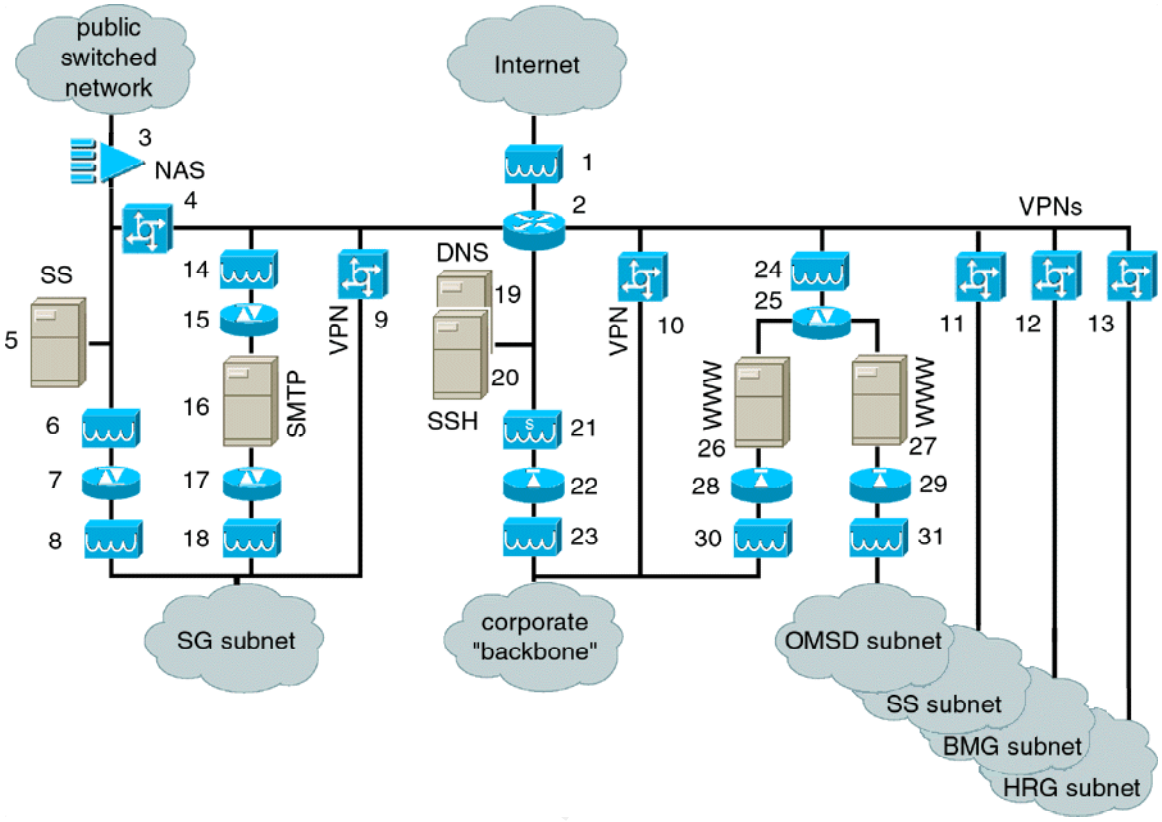

Figure 6. GIAC Enterprises security gateway design

1.5.2.1 Gateway element descriptions (including application specifications)

border router

Element 1—inbound packet filter used to block known bad addresses (e.g., private or spoofed addresses) and unwanted protocols as early as possible.

Element 2—multi-port router used to segregate and log incoming traffic.

remote access gateway

Elements 3 and 4—network access servers (NAS) that provide connectivity to restricted services for remote dialup and Internet users through the use of two-factor authentication.

Element 5—security server (SS) used in support of NAS to handle the processing of token challenge and response protocols. For the purposes of this paper, it is assumed that the token card used by remote users will be a SecurID, and thus the SS will have to be a RSA ACE/Server.[13] (Although it should be technically possible to use SmartCards and open source software to perform this function, I am not aware of a working, multi-platform solution to this problem.)

Element 6—inbound packet filter used to screen out protocols not supported by the element 7 proxy.

Element 7—a reverse http/https proxy (Apache) to support access to corporate services. Includes Webmail application for email retrieval. (Note that both applications are open source.)

Element 8—packet filter used to limit connectivity with element 7 to the internal corporate servers.

VPN gateway

Elements 9, 10, 11, 12, and 13—VPN elements used in a bridging role to cross connect various corporate department and group subnets that are physically located at different sites.

SMTP email gateway

Elements 14 and 18—packet filters used to screen out non-SMTP protocols (both) and, for element 18, limit connectivity with element 17 to the internal mail server(s).

Elements 15 and 17—email proxies (wrappers). The inbound (15) proxy is envisioned to be *smwrap*, part of the T.Rex firewall[14] kit, since it provides protection against multiple types of SMTP-based attacks. The outbound element (17) only has to be a simple proxy to provide the desired isolation between the gateway

[13] http://www.rsasecurity.com
[14] http://www.opensourcefirewall.com/trex_functional_summary_body.html

and the internal network; the TIS toolkit[15] components *smap* and *smapd*, or the Juniper toolkit[16] components *smtpd* and *smtpfwdd,* are suitable for this role. (Note that these applications are open source.)

Element 16—email content analyzer. Although it is possible to screen for undesired subject material, the concern here is to eliminate the threat posed by viruses. Presumably due to the intense, continuing effort involved in developing and maintaining the necessary viral signatures for an effective tool, little is available in terms of an open source email virus scanner. Rather, here it is suggested a product be used from the leader in the field (based on market share): the McAfee Webshield[17] from Network Associates running on their e50 appliance.

Domain name server

Element 19—The GIAC Enterprises maintains a split DNS in order to avoid unnecessary exposure of internal network details. The DNS server (and secondary) in the gateway only provides the necessary name resolution (and MX records, etc.) for external clients to connect to the front end of the gateway (not even all gateway devices are listed). This DNS sever will also be used for chaining off internal requests (lookups) to other servers. The DNS will be based on the latest version of BIND (open source).

Secure shell server

Element 20—As discussed in section 1.2.1, contract buyers will retrieve their purchases via *ssh* (*scp*; although this could also be ftp tunneled back over ssh). After contract agreement is reached, materials and an appropriate certificate (for authentication) will be uploaded to this server for some specified time (some few number of days or until downloaded, whichever comes first) in order to minimize exposure to potential compromise threats.

Internet access gateway

Element 21—dynamic packet filter. While the other gateways (e.g., remote access and SMTP) have additional layers of protection before the "backend" firewall is reached, that is not the case for the Internet access set of proxies. However, since this gateway only provides a path for internal clients to access web-based services, a dynamic (or stateful inspection) filter can be used to provide additional protection.

Element 22—proxy server. To support the identified outgoing connectivity requirements, application and

[15] http://www.tis.com/research/software/fwtk_over.html
[16] http://www.obtuse.com/juniper
[17] http://www.mcafeeb2b.com

circuit-level proxies are installed and used, as appropriate, to support the following protocols: PASV ftp, ssh, http/https, and DNS. All of the necessary proxies can be found open source (including the TIS, Juniper, and T.Rex firewall toolkits (see above), Squid,[18] and SOCKS[19]). Certain services that are potentially more risky (i.e., ftp) can also be limited to authorized users (e.g., members of the SG) by proper proxy selection (e.g., for ftp, the T.Rex ftp proxy includes access controls).

Element 23—packet filter. Can be used to help restrict access to firewall services to authorized machines (e.g., access to the ssh server is limited to machines within the CMSD).

WWW gateway

Element 24—inbound packet filter used to screen out protocols not supported by the element 25 proxy.

Element 25—a reverse http/https proxy (Apache) running on a multi-port "box" to support protected access to open corporate information, including that posted by the CWD and CMSD, as well as to the e-commerce server(s) of the OMSD.

Elements 26 and 27—Apache web servers.

Elements 28 and 29—ssh proxies.

Elements 30 and 31—packet filters. Both are used to screen out non-ssh protocols. Element 30 can also be used to help restrict connections to authorized machines. Note the physical connectivity to these filters helps segregate and protect the sensitive, OMSD data.

other considerations

1. All gateway machines should be treated as bastion hosts and locked down as much as possible.

2. As in the internal network description of section 1.5.1, NID should be used to instrument and verify the proper performance of network security components. While the internal network is responsible for monitoring the security gateway performance, NID should be deployed to monitor the forward elements inside the gateway (e.g., the 14/15 and 24/25 firewalls).

[18] http://squid.nlanr.net/Squid/
[19] http://www.socks.nec.com

3. All gateway machines should be located in an exclusive, physically secure room. Because of the relatively few number of machines, management, backup, and "remote" logging can all be accomplished out-of-band locally (e.g., no requirements to proxy SNMP through the firewalls).

1.5.2.2 Gateway vendor selections (including OS considerations)

router— In terms of the brand and version of border router to use: On the basis of safeguard selection principles already discussed (Parker), it would be nice to use Linux-based routers.[20] However, the scarcity of performance and supported configuration data makes this option suspect without much more research and possibly testing. Rather, the Cisco line of routers was selected because of their experience base and apparent viability as a company. Based, then, on the size and expected growth of GIAC Enterprises, the router of choice[21] is the Cisco 3660 router. The only concern here is that, at least for the present, modules are not available for this series that support gigabit Ethernet; if further analysis of connectivity requirements indicates that short-terms needs will quickly demand this level of performance, then a Cisco 7500 router should be used. (Also note that no effort has been made herein to evaluate the possibility of providing a multiservice—data, voice and video traffic—network.)

NAS— For uniformity or compatibility sake (e.g., an OS (Cisco IOS), and possibly some hardware modules, in common with the router), Cisco devices were selected to provide access services; another consideration was the support for RADIUS (used by the security server). For dialup support, the Cisco AS5300 Universal Access Server was selected, while a 2600-series router with two Ethernet ports will work fine for Internet access (most any Cisco router will work).

backend firewalls—Linux based, multi-homed servers containing the filters and proxies connected to the internal networks. An easy solution might be to purchase T.Rex appliances. (This configuration will also be used for internal firewalls.)

frontend firewalls—OpenBSD-based, multi-homed servers containing the filters and proxies connected to the border router. OpenBSD was selected for this role over Linux in order to try to avoid common modes of failure.

email content analyzer and security server—Sparc with Solaris OS. The selected software is available to run under Solaris; while this is not open source, excellent procedures (such as from SANS) are available for securing this platform.

---

[20] http://www.linuxrouter.org/
[21] Cisco Systems, *Multi Service Access Solutions*

other servers—Linux is the open OS of choice due to the high-level of active development and support.

VPNs—As for the router, it would be nice to select an open source platform like Linux, but it is not felt that the maturity is there. Rather, the Red Creek[22] Ravlin 7160 was selected for use with the IPsec protocol. A non-integrated, firewall bypass configuration was selected in order to avoid being tied to any particular vendor solution and protocol issues.

2.0 Security Policy

Note: The policies discussed herein are in support of the architecture and data flows discussed in section 1. All rule sets implementing security policy will be tested in a lab setting using packet creation tools like NMAP on the front end and TCPdump on the backend.

2.1 Problem Statement

assignment 2 - Security Policy (25 Points)
Based on the security architecture that you defined in Assignment 1, provide a security policy for AT LEAST the following three components:

Border Router
Primary Firewall
VPN

You may also wish to include one or more internal firewalls used to implement defense in depth or to separate business functions.
By 'security policy' we mean the specific ACLs, firewall ruleset, IPSec policy, etc. (as appropriate) for the specific component used in your architecture. For each component, be sure to consider internal business operations, customers, suppliers and partners. Keep in mind you are an E-Business with customers, suppliers, and partners - you MAY NOT simply block everything!
(Special note VPNs: since IPSec VPNs are still a bit flaky when it comes to implementation, that component will be graded more loosely than the border router and primary firewall. However, be sure to define whether split-horizon is implemented, key exchange parameters, the choice of AH or ESP and why. PPP-based VPNs are also fully acceptable as long as they are well defined.)
For each security policy, write a tutorial on how to implement each ACL, rule, or policy measure on your specific component. Please use screen shots, network traffic traces, firewall log information, and/or URLs to find further information as appropriate. Be certain to include the following:

1. *The service or protocol addressed by the ACL or rule, and the reason these services might be considered a vulnerability.*
2. *Any relevant information about the behavior of the service or protocol on the network.*
3. *The syntax of the ACL, filter, rule, etc.*
4. *A description of each of the parts of the filter.*
5. *An explanation of how to apply the filter.*
6. *If the filter is order-dependent, list any rules that should precede and/or follow this filter, and why this order is important. (Note: instead of explaining order dependencies for each individual rule, you may wish to create a separate section of your practical that describes the order in which ALL of the rules should be applied, and why.)*
7. *Explain how to test the ACL/filter/rule.*

Be certain to point out any tips, tricks, or "gotchas".

[22] http://www.redcreek.com

2.2 Border Router

The policy set forth for the Cisco border router below is presented in the form of a commented rule set that could be copied and pasted into a terminal window attached to the router as a virtual console. (Filter rule sets can be added to, but not edited; they have to be deleted and completely reentered line by line with the desired changes unless this "trick" is used!) This policy does not cover hardening the router itself.[23]

```
!Cisco packet-filtering router access control lists (ACLs)
!extended ACLs used in order to filter on header data in addition
!to IP addresses available under standard ACLs
!
!notes:
!(1) depending upon the actual addressing scheme used
!any rule below with <internal address> may, in fact, require
!multiple instances for each separate address range
!(2) since the rules are tested in order for a match,
! the "logical" order presented below is unlikely to be the best for performance.
!Use a command like "show running" to find out the number of hits each rule
!receives for some time period, and adjust as desired.
!
!beginning of access-list 101
!used for inbound (ingress) traffic through the external interface of the border router
!
!!!!!!!!!! bad address concerns !!!!!!!!!!
!
!deny private addressing per RFC 1918
access-list 101 deny ip 192.168.0.0 0.0.255.255 any
access-list 101 deny ip 172.16.0.0 0.15.255.255 any
access-list 101 deny ip 10.0.0.0 0.255.255.255 any
!deny localhost, broadcast, and multicast addresses
access-list 101 deny ip 127.0.0.0 0.255.255.255 any
access-list 101 deny ip 255.0.0.0 0.255.255.255 any
access-list 101 deny ip 224.0.0.0 7.255.255.255 any
!deny packets with NULL IP address
access-list 101 deny ip host 0.0.0.0 any
!deny spoofed (internal) addresses
access-list 101 deny ip <internal address> <internal mask> any
!deny spoofed router external interface address
access-list 101 deny ip host <router external interface IP address> any
!
!!!!!!!!!! vulnerable tcp-based service concerns !!!!!!!!!!
!
!since we are about to allow any established TCP connections (simple test of ACK bit)
!for ports >1023 we have to be aware of high-port services that might be in use internally,
!or common exploit ports,
!that we don't want an attacker to be able to "hit" or exploit them (as a set ACK bit is easy to fake)
!Note that this list should be reviewed periodically and updated as appropriate.
!SOCKS
```

[23] E.g., see http://www.cisco.com/warp/public/707/21.html

```
access-list 101 deny tcp any any eq 1080
!Citrix ICA
access-list 101 deny tcp any any eq 1494
!NFS
access-list 101 deny tcp any any eq 2049
!lockd
access-list 101 deny tcp any any eq 4045
!SGI object server
access-list 101 deny tcp any any eq 5135
!x-windows
!although note that this list may have to be extended to cover the range of 6000 through 6255
access-list 101 deny tcp any any eq 6000
access-list 101 deny tcp any any eq 6001
!Ircd
access-list 101 deny tcp any any eq 6667
!common high-port HTTP
access-list 101 deny tcp any any eq 8000
access-list 101 deny tcp any any eq 8080
access-list 101 deny tcp any any eq 8888
!
!!!!!!!!!! allowed tcp-based connections !!!!!!!!!!
!
!allow tcp-based “established” connections
!as this is the most likely route for attackers to probe or compromise our network, log the connections
access-list 101 permit tcp any <internal address> <internal mask> gt 1023 established log
! allow http and https connections to web servers
!as we are interested in sensitive data retrieved via SSL, log the connections
access-list 101 permit tcp any gt 1023 host <WWW gateway external IP address> eq 80
access-list 101 permit tcp any gt 1023 host <WWW gateway external IP address> eq 443 log
! allow smtp connections to mail gateway
access-list 101 permit tcp any host <mail gateway external IP address> eq smtp
! allow ssh connections to external ssh server
!as we are interested in sensitive data retrieved via SSH, log the connections
access-list 101 permit tcp any host <ssh server IP address> eq ssh log
! allow IPsec protocol connections between main office and remote location VPN devices
access-list 101 permit 50 host <VPN gateway “mate” IP address> host <VPN gateway external IP address>
access-list 101 permit 51 host <VPN gateway “mate” IP address> host <VPN gateway external IP address>
!
!!!!!!!!!! vulnerable udp-based service concerns !!!!!!!!!!
!
!since we are about to allow UDP connections in response to DNS queries,
!for ports >1023 we have to be aware of high-port services that might be in use internally,
!or common exploit ports, that we don’t want an attacker to be able to “hit” or exploit them
!Note that this list should be reviewed periodically and updated as appropriate.
!NFS
access-list 101 deny udp any any eq 2049
!lockd
access-list 101 deny udp any any eq 4045
!Ircd
access-list 101 deny udp any any eq 6667
!NetBus
access-list 101 deny udp any any eq 12345
access-list 101 deny udp any any eq 12346
access-list 101 deny udp any any eq 20034
```

```
!Back Orifice
access-list 101 deny udp any any eq 31337
!
!!!!!!!!!! allowed udp traffic !!!!!!!!!!
!
! allow udp responses to DNS queries
access-list 101 permit udp any eq 53 <internal address> <internal mask> gt 1023
!allow queries to our external DNS (note: no tcp (zone) transfers enabled)
access-list 101 permit udp any gt 1023 host < external primary DNS IP address> eq 53
access-list 101 permit udp any gt 1023 host < external secondary DNS IP address> eq 53
! allow IPsec connection setup between main office and remote location VPN devices
access-list 101 permit udp host <VPN gateway "mate" IP address> host <VPN gateway external IP address> eq
389
access-list 101 permit udp host <VPN gateway "mate" IP address> host <VPN gateway external IP address> eq
500
!
!!!!!!!!!! icmp-based service concerns !!!!!!!!!!
!
!some people prefer to "deny icmp any any" but they can be useful! duh!
!here take the stance of blocking icmp that can be used for "probing" by an attacker
!HOWEVER, this will hinder us from troubleshooting main office-to-remote office problems!
access-list 101 deny icmp any <internal address> <internal mask> echo-request
!
!!!!!!!!!! icmp we would like to get back based on responses to our own outgoing traffic !!!!!!!!!!
!
access-list 101 permit icmp any <internal address> <internal mask> net-unreachable
access-list 101 permit icmp any <internal address> <internal mask> host-unreachable
access-list 101 permit icmp any <internal address> <internal mask> port-unreachable
access-list 101 permit icmp any <internal address> <internal mask> packet-too-big
access-list 101 permit icmp any <internal address> <internal mask> administratively-prohibited
access-list 101 permit icmp any <internal address> <internal mask> source-quench
access-list 101 permit icmp any <internal address> <internal mask> ttl-exceeded
access-list 101 permit icmp any <internal address> <internal mask> echo-reply
!
!anything left, deny (the default), but here call out in order to log at least initially for trouble shooting purposes
access-list 101 deny ip any any log
!
!end of access-list 101
!
!beginning of access-list 102
!used for outbound (egress) traffic through the external interface of the border router
!
!!!!!!!!!! icmp-based service concerns !!!!!!!!!!
!
!here take the stance of blocking icmp that might be a return from "probing" by an attacker
!denying much of the information we would like!
!HOWEVER, this will hinder us from troubleshooting main office-to-remote office problems!
access-list 101 deny icmp any any net-unreachable
access-list 101 deny icmp any any host-unreachable
access-list 101 deny icmp any any port-unreachable
access-list 101 deny icmp any any administratively-prohibited
access-list 101 deny icmp any any ttl-exceeded
access-list 101 deny icmp any any echo-reply
!
```

```
!!!!!!!!!! anything else goes from valid internal addresses !!!!!!!!!!
!
access-list 102 permit ip <internal address> <internal mask> any
access-list 102 permit ip <internal address> <internal mask> any
!
!!!!!!!!!! deny invalid internal addresses and log !!!!!!!!!!
!
access-list 102 deny   ip any any log
!
!end of access-list 102
```

2.3 Primary Firewall

For the purposes of this paper, the primary firewall is considered to be the Linux-based one providing outgoing connectivity (element 21 in Figure 6). The policies for this firewall will be implemented using the stateful packet filtering available with *netfilter*[24] (*iptables* provides user command interface) under Linux 2.4. The policies do not include hardening the firewall itself.[25] Commands to add, delete or insert rules can be interactive (one at a time like router ACL commands discussed above—although the delete and insert is much nicer!), or preferably, they can be collected into a script file that is executed during the system boot process (nicer still!). By switching to a stateful filter, especially in a case where all connections originate on one side of the interface as here, the rule set is greatly simplified:

```
# Allow all packets out of the internal network
iptables -A FORWARD -m state --state NEW –i $INT_IFACE -s $OUR_NETWORK -j ACCEPT
# and allow the packets associated with those connections back in.
iptables -A FORWARD -m state –state ESTABLISHED,RELATED -i $EXT_IFACE -s ! $OUR_NETWORK -j ACCEPT
#drop all other packets
iptables –A FORWARD -i $EXT_IFACE –j DROP
```

And that's it! Even if the filters in the border router fail, this stateful filter will continue to provide protection from undesired packets to the proxies (element 22).

2.4 VPN

VPN connectivity between the main and remote offices will be established using the Red Creek Ravlin 7160. Setup is very straight forward using the front panel and only a few steps (like establishing the local and remote IP addresses and masks). In this application, the 7160s will be configured to act like a bridge (remote and local networks within any one department or group will share a common subnet address space). All connections will be tunneled using the Encapsulating Security Payload (ESP). Authentication will be based

[24] http://netfilter.kernelnotes.org contains the necessary references for configuring the Linux kernel to use the features available with netfilter.
[25] E.g., the guides available at www.sans.org

on Hashed Message Authentication Codes (HMAC) using MD5. Encryption will be based on 168-bit triple DES.

3.0 Audit

3.1 Problem Statement

Assignment 3 - Audit Your Security Architecture (25 Points)
You have been assigned to provide technical support for a comprehensive information systems audit for GIAC Enterprises. You are required to audit the Primary Firewall described in Assignments 1 and 2. Your assignment is to:

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

2.Implement the assessment. Validate that the Primary Firewall is actually implementing the security policy. Be certain to state exactly how you do this, including the tools and commands used. Include screen shots in your report if possible.

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

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

3.2 Planning

1. Establish a clear, concise, signed statement of work (contract) that defines the security component(s) to be tested (here it is the primary firewall represented by element 21 of Figure 6) and the threat agent(s) to be considered (here assume an external, cyber-based threat only; could be internal, external with internal assistance, physical presence, etc.). Planning costs should be provided (say the equivalent of two man days).

2. Gather and study relevant documentation. At a minimum this should include any rule sets implemented. Ideally this would also include: corporate security policies; hardware, OS and application configuration procedures and records; security test plans and records; and recent logs (firewall and NIDs).

3. Search for possible exploits and vulnerabilities for hardware/OS/application in use and any relevant updates or patches.

4. Develop a profile for the security component of expected responses to various stimuli. Identify appropriate test equipment to generate these stimuli and measure responses, and the time necessary to conduct the tests, analyze the test data, and produce a final report (say one, one, and two days respectively). Identify any "out of band" test parameters that should be considered that are not reflected in the existing rule sets or policies.

5. Conduct a meeting between the parties concerned (e.g., the appropriate manager(s), security personnel, and operators/administrators). Review the findings to date and test plans. Document understandings developed, including modifications to the statement of work, if required. Have a clear understanding of what personnel will be providing test support (including implementation of a recovery plan, if required). Make sure full backups have been made (and that restores actually work!).

3.3 Implementation

1. After the plans have been finalized and duly approved, they can be implemented.

2. For the present case, testing the primary firewall can best be done out-of-hours and out-of-loop. This is a result of the gateway design—the primary firewall only provides Internet connectivity for internal clients, and GIAC Enterprises (it is here assumed) is a one-shift operation. WWW servers will remain unaffected.

3. Once brought off line, the firewall rule sets should be reloaded with full logging turned on. NIDs with TCPdump running should be connected so as to fully instrument both sides of the firewall. When this has been completed and event logging has been verified, the tester can begin applying the desired stimuli to the external interface; nmap[26] is a useful tool for this purpose (although a number of other network-based vulnerability scanners are available that could be used).

4. Compare the test logs against the expected responses, noting any discrepancies.

5. Restore the test item to its original configuration and verify its proper operation with site personnel.

6. Document the results, and distribute as defined in the statement of work.

3.4 Perimeter analysis

1. While the stateful filter appears to be a significant improvement in managing packet traffic, it is not at all clear that data is available on how robust the filter itself is nor on how immune the traffic is to hijacking. And while the rule set as implemented above did not duplicate the rules found on the border router, consideration should be made for doing so. Or, if one of the DNS or SSH servers is compromised, are there additional firewall rules that should be considered? (Part of the problem with developing or assessing any

[26] http://www.insecure.org/nmap/index.html

security architecture design is the lack of any qualitative or quantitative measures that can be used in the decision making process.)

2. While the assumed focus of the assessment was on external threats, internal threats are a fact of life. With this in mind, the existing firewall design (or gateway for that matter) provides a high-bandwidth "leakage" path for internal information. Although it may mean loss of service, consideration should be given to additional restrictions in the egress rule set. (But it is not clear that any service can be provided that cannot be turned into a tunnel.)

3. The proxies located behind the firewall (in particular, those used for the http/https, ssh, and PASV ftp) do little or nothing to protect internal systems from viruses or other malware. A more robust design would implement gateways more like that specified above for email.

4.0 Test

4.1 Problem Statement

Assignment 4 - Design Under Fire (25 Points)

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

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

1.An attack against the firewall itself. Research vulnerabilities that have been found for the type of firewall chosen for the design. Choose an attack and explain the results of running that attack against the firewall.
2.A denial of service attack. Subject the design to a theoretical attack from 50 compromised cable modem/DSL systems using TCP SYN, UDP, or ICMP floods. Describe the countermeasures that can be put into place to mitigate the attack that you chose.
3.An attack plan to compromise an internal system through the perimeter system. Select a target, explain your reasons for choosing that target, and describe the process to compromise the target.

Note: this is the second time this assignment has been used. The first time, a number of students came up with magical "hand-waving "attacks. You must supply documentation (preferably a URL) for any vulnerability you use in your attack, and the exploit code that you use to accomplish the attack. The purpose of this exercise is for the student to clearly demonstrate they understand that firewall and perimeter systems are not magic "silver bullets" immune to all attacks.

4.2 Network design (http://www.sans.org/y2k/practical/Don_Munroe_gcfw.doc)

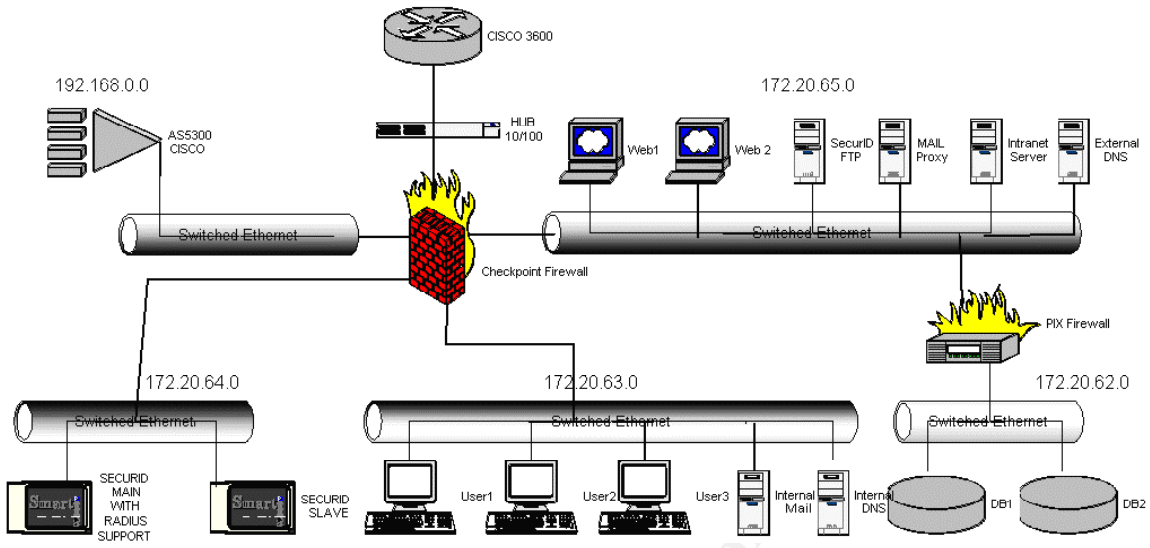

4.3 Firewall Vulnerabilities

A web search readily yields a number of vulnerabilities for the Checkpoint firewall. The most extensive list[27] was published at the Black Hat Briefings 2000, along with source code for several of the attacks. However, perhaps the simplest attack, a denial of service (DOS), was originally documented by Lance Spitzner[28] as part of his research[29] into how FW-1 handles fragmented packets.

Vulnerability—FW-1 does not inspect, nor does it log, fragmented packets until the packet has first been completely reassembled. Since these exploit packets are never fully assembled (they are placed in memory waiting for the rest of the packet to arrive), they are never inspected nor logged. Thus, the firewall rule base cannot be used to protect against the attack.

Exploit— Most fragment-based attacks that use incomplete or illegal fragments will work, including jolt2.[30]

[27] http://www.dataprotect.com/bh2000/blackhat-fw1.html
[28] http://packetstorm.securify.com/0006-exploits/firewall-1.fragment.txt
[29] http://www.enteract.com/~lspitz/fwtable.html which includes several vulnerabilities.
[30] http://securityfocus.com/data/vulnerabilities/exploits/jolt2.c

It is suggested that a few hundred fragments are sufficient to execute this DOS. The firewall does not have to be attacked directly, if the fragments are routed through the firewall for a system behind the firewall, FW-1 is still taken out.

4.4 Distributed DOS

> "The truth is that a site cannot defend itself from DDoS attacks alone. DDoS attacks depend upon the 'community' of the Internet and defenses, therefore, depend upon the Internet community acting like a community with a common interest. And defending against attack includes ensuring that our own sites are not the source of attacks and our own networks do not forward attacks."[31]

Based on CERT Advisory CA-1999-17,[32] the following actions should be taken to reduce the risk from a distributed DOS (DDOS) attack:

1. Implement network ingress filtering in accordance with RFC-2267.[33] In summary:
   - perform egress filtering (only allow valid internal addressing out)
   - block incoming broadcast addresses
   - turn off directed broadcast capability
   - block private and reserved addresses
   - block unused ports and those known to be associated with DDOS attacks
2. Implement the "Suggestions for System Administrators" found in the *Results of the Distributed-Systems Intruder Tools Workshop* sponsored by the CERT Coordination Center in December, 1999.[34]
3. Install the Open Transport 2.6 upgrade for any Macintosh computers on the network.[35]
4. Use personal firewalls on all systems.
5. Have Zombie Zapper[36] available to halt a zombie from flooding.
6. Consider use of dual-homed, fail-over Internet connections.
7. And, although not specific to DDOS protection, follow CERT, SANS, and TruSecure (ICSA) best practices procedures (e.g., keep current with software updates, perform intrusion detection, etc.).

4.5 Attack Plan

[31] http://www.sans.org/infosecFAQ/threats/DDoS.htm
[32] http://www.cert.org/advisories/CA-1999-17.html
[33] http://info.internet.isi.edu:80/in-notes/rfc/files/rfc2267.txt
[34] http://www.cert.org/reports/dsit_workshop.pdf
[35] http://www.csc.gatech.edu/macattack/ and http://www.csc.gatech.edu/macattack/macattack.html
[36] http://packetstorm.securify.com/distributed/zombie/

The first target of choice is one of the "user" systems located on the internal switched Ethernet LAN. The reasoning behind this choice is that if a seat is gained on the internal network, it will be much easier to then stage an attack against the ultimate targets: the databases residing behind the Cisco PIX Firewall.

Step 1—Conduct network scans in order to identify the OS type and version on the user platforms, commonly called fingerprinting,[37] by using tools such as *nmap*.

Step 2—As performed for the firewall in section 4.3, use the results of Step 1 to conduct a (web) search for user-host vulnerabilities.

Step 3—Compare the firewall and user-host vulnerabilities in order to identify one or more exploits that have a potential to succeed. Note that the existing vulnerabilities and associated exploit requirements may require that the border router be fingerprinted and compromised as well.

Step 4—Conduct the attack.

Step 5—Elude network- and host-based intrusion detection systems[38] and install "backdoors" on the user-host for later, easy access and use.

[37] E.g., http://www.insecure.org/nmap/nmap-fingerprinting-article.html see also http://packetstorm.securify.com/papers/protocols/host-detection.txt
[38] E.g., http://www.iwar.org.uk/comsec/resources/honey-pod/honeynet_papers/honeypot/ids.pdf