Global Information Assurance Certification Paper

Eric Peeters

April 15, 2005

GCFW Practical

Version 4.1

Assignment 1

Abstract

When wireless networks first came out, they were deployed as an extension of the wired network with little thought for security. If a company turned on a wireless network, it was performing the equivalent of hanging a network cable out in the street for anyone to plug any device into the internal network.

The first measure of security for wireless network was Wired Equivalent Privacy (WEP), which uses 64-bit, 128-bit and sometimes 156-bit encryption keys to create an encrypted tunnel between a wireless client and an access point. Unfortunately, a flaw in WEP that rendered it very easy to break was quickly discovered[1].

Next came 801.1x Wifi Protected Access (WPA) which is itself a precursor to 802.11i. WPA's contributions to wireless network security include improved encryption as well as cross-authentication of the client and access point using 802.1x EAP.

The 802.11i standard has recently been approved by the IEEE but its adoption by professional users is slow at best, due to concerns over the cost and complexities associated with rolling out new services and acquiring new hardware. 802.11i brings to wireless security yet another new set of encryption algorithms as well as an improved cipher, Advanced Encryption Standard (AES), and faster roaming between access points.

Today, most businesses are aware of the risks associated with a wireless network and treat them with as much caution as they do inbound traffic from the internet, but there are still some than can be found with no security at all. Wireless networks are a great convenience and a productivity boosting tool for many companies and employees, and they are therefore here to stay.

By their very nature, GIAC Enterprises' scanners and laptops have different connectivity needs and different capabilities. They should therefore have their dedicated wireless network to ensure that whatever security is added to this network is the best fit.

[1] Borisov Nikita, Goldberg Ian, Wagner David. Security of the WEP algorithm. http://www.isaac.cs.berkeley.edu/isaac/wep-faq.html

Scanners

If GIAC Enterprises uses scanners in its shipping department, it is safe to assume their purpose is to record product movements or display their status. Both functions require little more than a connection to an inventory database, and this is exactly the kind of access the network diagram below will provide.

Wireless Scanner Diagram

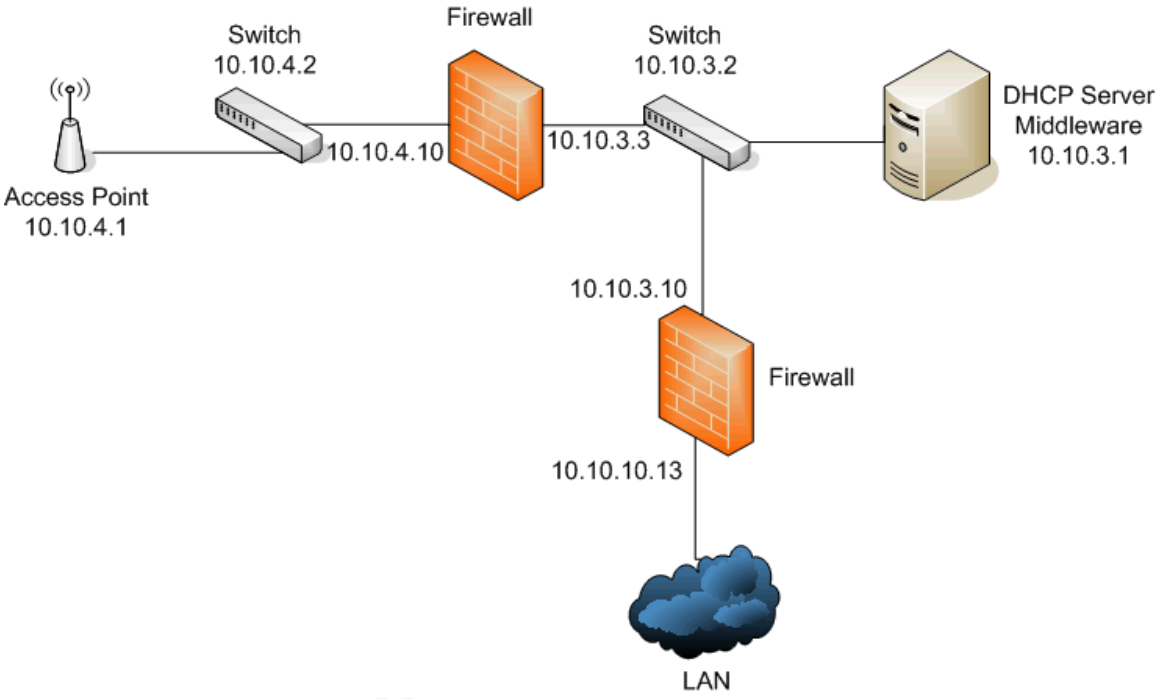

WEP vs. WPA

Wireless technology has made tremendous progress when its comes to authentication over the past few years, but it is the author's professional experience that this progress is very slow in trickling down to handheld scanners. While brand-new scanners do support WPA and even claim to be 802.11i ready, there are many older-design scanners still available on the marketplace where the only wireless security protocol available is WEP. A key feature of WPA is the need for the client and the access point to authenticate each other using one of several 802.1x Extensible Authentication Protocol (EAP) methods.
Even with newer scanners, however, WEP is often preferred to WPA for a number of reasons:

- Despite Cisco's dominant market share in wireless scanner radios, its proprietary EAP method, EAP-LEAP (Lightweight EAP) has fallen in

disrepute since the release in October 2003 of Asleap in October 2003, a very efficient brute force password cracker for LEAP. Cisco's newer EAP method, EAP-FAST, was released to counter Asleap, but has had a limited following so far.

- EAP-TLS is the most secure EAP method, and the easiest to use from the end-user's point of view because it requires only that the client and access point exchange their respective authentication certificate. The need for such a certificate for each device makes it difficult to manage, not to account for the fact that many basic scanners were never designed to hold their own certificate.
- EAP-TTLS is similar to EAP-TLS but easier to deploy because the client authenticates itself using a user name and password combination sent over an encrypted link rather than a certificate. While this method addresses the problems associated with issuing and storing in each device a certificate, EAP-TTLS isn't compatible with many shop floor scanners that are not designed to allow users to enter a user name and password.
- EAP-PEAP supports client authentication either through a user name and password combination, or through a token, either method bumping against issues already raised with EAP-TLS and EAP-TTLS.

While WPA would undoubtedly give a better level of security to GiAC Enterprises' wireless scanner network, we will therefore assume that the security of choice is WEP and build a security infrastructure around it.

Access Point

Any WEP-enabled access point that can operate as a closed system (turn off ESSID broadcasting) and perform MAC address filtering will be sufficient for GIAC Enterprises' needs. Depending on the size of the shipping department, GIAC Enterprises might consider the use of so-called thin access points. Thin access points, as opposed to fat access points, do not perform any processing other than the one associated with the radio signal itself. Thin access points are connected to a common wireless switch where WEP negotiation and MAC address filtering will take place. The main advantage of a wireless switch over fat access points is that rolling out additional access points require little more than connecting them to the wireless switch. On the other hand, wireless switches are still fairly expensive and probably not worth the extra cost if only a few access points are required to provide adequate coverage.

Cisco, with its Aironet 1200 Series access points[2], is the largest vendor of fat access points, but the author prefers the Orinoco line from Proxim[3] for its breadth of options and lower price point. In the thin access point market,

[2] http://www.cisco.com/en/US/products/hw/wireless/ps430/index.html
[3] http://www.proxim.com/products/wifi/ap/

Extreme Networks[4] is one of the most reputed vendors so far.

Server

The presence of a server in our wireless network has a dual purpose. The server will issue IP addresses to all scanners through a DHCP server, but the server will also act as middleware between the scanner and the SQL database.

Some argue that all addresses in the DHCP scope should be reserved to a specific MAC address corresponding to one of the scanners in use on the floor. It is the author's opinion that this step adds a level of complexity without any pay-off. If a malicious user can change his device's MAC address to that of a device with an authorized address in the access point's MAC table, this device will automatically be recognized by the DHCP server and issued its corresponding IP address. A 10.10.4.128/25 address block should be enough for the DHCP scope.

Many mobile operating systems (including all Windows Mobile versions, except the most recent one, Windows Mobile 2003) are unable to communicate directly with a SQL server, making a middleware application necessary. Even if such an application isn't required, it still should be used because it can play a large role in securing a network by adding a layer between the SQL sever, usually located on the internal network, and the scanner, so that even if a malicious user were to gain unauthorized access to the wireless network, he or she would still have no direct path to the internal network.

Odyssey Software's CEfusion[5] is a good example of a middleware application intervening in the communication between a SQL server and a scanner. An agent on the scanner redirects all database interactions to the port CEfusion is listening to on the server (TCP/25250 by default) and forwards them to the SQL server via ODBC. CEfusion has the additional advantage that it does not need to run under a system account, so even if a hacker were to be able to use an injection attack against CEfusion, he or she would gain only the access granted whatever account CEfusion is running under rather than complete access to the server.

External Firewall

It has long been established that WEP isn't very hard to break, and any semi-experienced wireless hacker will have no problem piercing the thin layers of extra security that MAC address filtering and the lack of ESSID broadcasting add. It is therefore strongly recommended to protect the middleware/DHCP

[4] http://www.extremenetworks.com/libraries/prodpdfs/products/UnifiedWireless.asp
[5] http://www.odysseysoftware.com/products_cefusion.asp

server with an external firewall that would allow through only incoming packets with a destination matching the server's IP and the ports used either by the DHCP server (UDP/67 and UDP/68) or by the CEfusion middleware application (TCP/25250). For obvious security reasons, the server will be in a different subnet than the scanners, so the external firewall will have to be configured to pass DHCP requests through since those are normally not routed. The Cisco PIX 515E may be configured as such, and is already used elsewhere on GIAC Enterprises' network (see section 2), so it would be a perfect candidate for the role of external firewall. Since all database interactions will be initiated by the scanner, the Cisco PIX firewall should be configured to block all outbound connections to the wireless network, except that which is necessary to manage the access point and the switch (generally TCP/443 for HTTPS management or TCP/23 for telnet management).

Internal Firewall

A Cisco PIX 515E may also be used as the internal firewall to separate our server subnet from the internal network. Since the CEfusion middleware application forwards all packets destined for the SQL server via ODBC, the only accepted inbound traffic should have as source IP the middleware server and as destination IP the SQL server's internal IP address, with TCP/1433 as the destination port. Allowed outbound traffic from the internal network should be limited to access point and switch management traffic (see above) as well as Remote Desktop Protocol (TCP/3389) destined for our Windows 2003 middleware server.

Laptops

Laptops have more connectivity requirements than scanners, for instance accessing the company Intranet, the SQL database or the mail server, but they should also be able to secure their wireless connection using WPA as opposed to WEP, since both Windows 2000 and Windows XP have built-in support for several EAP methods.

Wireless Laptop Network Diagram

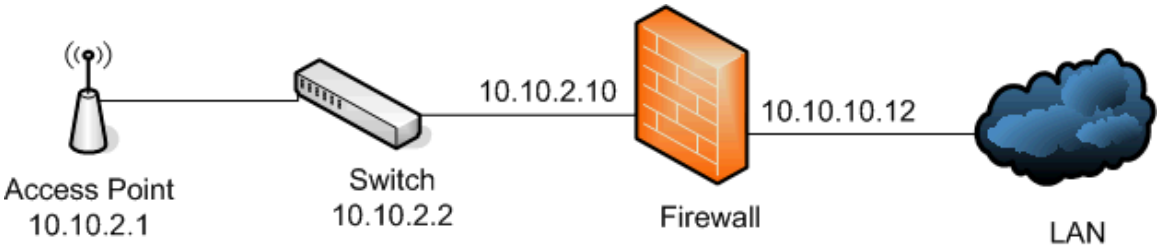

Access Point

The configuration of the wireless laptop network access point differs little from that of its scanner network counterpart. This access point should also be configured as a closed system and should use MAC filtering. Again, there is a choice to be made between thin or fat access points depending on the number to be deployed and the complexity of their configuration.

The key difference is that the laptop network access point will use WPA instead of WEP with EAP-PEAP as its 802.1x EAP method. Since support for EAP-PEAP is built in Windows 2000 and XP as discussed above, laptop users may not even have to re-enter their network logon to authenticate to the access point as cached credentials, if present, will be sufficient.

The access point also needs to be configured to use RADIUS to verify the network logon presented by the wireless laptop against a domain controller on the internal network. Since WPA is much more reliable than WEP, there is no need for a middleware server and two firewalls. In the absence of a server, though, the access point will have to do double-duty as DHCP server, a role that most enterprise-level access points can perform adequately. The address block 10.10.2.128/25 may be assigned to the DHCP scope.

Firewall

The firewall needs to allow several types of inbound packets into the internal network. First and foremost is RADIUS traffic. RADIUS listens on TCP/1812 by default, but it can be easily changed, so extra care should go in making sure that the firewall rule contains the appropriate destination port. The destination IP should obviously be that of the RADIUS server on the internal network. The source IP for this particular rule should be that of the access point.

Other rules should given access to the services required by the wireless port users, which at a minimum should be the company intranet (TCP/80) and probably the SQL server (TCP/1433), with the appropriate internal IP addresses as destination IPs. Service access rules should all specify as source the same IP range that is defined in the access point's DHCP scope. As with the wireless scanner network, the only outbound rule required is the one giving management access to the access point and the switch from the internal network.

Again, the Cisco PIX 515E would be a perfect fit as firewall, if only because it is already used elsewhere on the GIAC Enterprises network.

Other Security Options

While often ignored, the choice and location of antennas can play a sizeable role in a wireless security policy. If wireless coverage is needed only inside the warehouse, why would GIAC Enterprises deploy antennas extending coverage across the parking lot as well?

Most antennas in place are so-called omni directional antennas, meaning that they radiate their coverage in a circle over the horizontal plane. While omni directional antennas may be placed in the middle of a warehouse, they should never be installed against external walls for instance, where patch antennas or yagi antennas would be a better fit.

A patch antenna spreads its radio wave in a 180 degree angle both along the vertical and horizontal planes. A patch antenna is thus an ideal fit along a long straight external wall, since it will concentrate its signal strength towards the building rather than spreading it evenly in all directions, including outside the building.

A yagi antenna is similar to a patch antenna, but radiates its signal at an adjustable angle inferior to 180 degrees, making it the ideal antenna for corner locations.

Many enterprise-class antennas come with a power limiter that allows their signal strength to be dialed down or up as necessary to extend coverage where necessary without extending it too far. It is obviously impossible to provide coverage everywhere it is required and nowhere it isn't, but combining the right antenna selection and placement with the right signal strength ought to bring a wireless network closer to this objective and thereby lessen the risk that an unknown attacker will attempt to penetrate GIAC Enterprises' wireless network from the coffee shop across the street.

Several vendors, such as BlueSocket[6] or AirDefense[7], offer wireless security appliances with features ranging from rogue access points detection to role-based access control.

A rogue access point is generally installed by a company employee without the authorization of the IT department. These access points present a grave threat to network security because they are usually home-user products with little to no security and are plugged in directly into the internal network, making their detection paramount.

Role-based access control grants a wireless network user access only to specific subnets and/or services based on his role as determined by the

[6] http://www.bluesocket.com/
[7] http://www.airdefense.net/

wireless security appliance. Factors such as how the user authenticated, whether he authenticated at all or which access point he is accessing the wireless network from can be included in determining the user's role. This is particularly useful in large wireless environments with numerous users falling in different categories, such as visitors, mobile employees, external vendors, contractors, etc….

Conclusions

By using a two-network approach for its wireless scanners and laptops, GIAC Enterprises will be able to secure each network according to its specific characteristics, rather than have to devise a single network to accommodate wireless clients with different capabilities and connectivity needs.

At first glance, the scanner network may seem better protected than the laptop network with its use of two firewalls and a middleware application to prevent direct connections between scanners and the corporate database, but this added complexity is made necessary by the fact that the wireless connection is secured by WEP, a very inferior standard than that which is used for the laptop network.

To a varying degree, both wireless networks illustrate the defense-in-depth concept that is an overriding concern in designing GIAC Enterprises' security infrastructure. With both networks, the first line of defense is the absence of ESSID broadcasting, making it impossible for wireless clients that have not been configured with the right ESSID to even connect to an access point, much less authenticate themselves. The second and third lines of defense are also similar, in that wireless clients on both networks must have their MAC address included in the access point's authorized address tables, and they both must then authenticate themselves to the access point and negotiate an encryption key.

Due to the superior standard of its encryption and authentication method, the next line of defense on the laptop network is also the last, as it is the firewall separating the wireless sub network from the internal network. Clients on the scanner network, on the other hand, must go through a firewall before reaching a middleware application that will connect with the internal network on their behalf, and even that application has another firewall to traverse, all in order to make up for the lower standard that WEP is.

Finally, the security of both wireless networks will benefit from a careful selection of antennas and their power settings, while the internal network as a whole would be better protected from wireless threats with the implementation of a rogue access point detector.

Assignment 2

GIAC Enterprises' network needs to meet the access requirements of five different user groups all the while maintaining the highest degree of network security. The five use groups are:

- Employees on the internal network;
- Mobile employees (sales force);
- Customers;
- Suppliers;
- Partners.

In order to ensure that these five groups' connectivity needs will be met without bottlenecks, GIAC Enterprises will be connected to the Internet through a T-3 dedicated line. Each device used on the network will therefore have to be able to deliver at a minimum 45 Mbps of throughput.

Mobile employees will need to access the same applications on the internal network that in-house employees do, for which a VPN server will be dedicated to their use. Suppliers and partners will interact with GIAC Enterprise through a web interface powered by a database. To maintain integrity and confidentiality of the data, suppliers and partners will not be able to access the database directly. Rather, suppliers will use a secured form on a web server to upload new fortunes, while partners will use another secured form to receive new fortunes to be translated and post their translations. Customers will have access to another web server to browse GIAC Enterprises' catalog and place their order. Other than for maintenances and upgrades performed by the IT staff, internal users will not need access to either web server, as they will be able to view all pertinent data on an internal database.

The access for the various groups may be summarized as follows:

Source	Destination	Port	Protocol	Purpose
Local users	Internet	> 1024	TCP/UDP	Internet browsing
Customers	Web server 1	80	TCP	Catalog browsing
Customers	Web server 1	443	TCP	Secure ordering
Suppliers	Web server 2	443	TCP	Fortunes upload
Partners	Web server 2	443	TCP	Fortunes upload & download
Sales force	VPN server	500	UDP	Negotiating SA for VPN tunnel
Sales force	VPN server	IP Protocol 50 (ESP)		Remote access to internal network
Everyone except local users	DNS server	53	UDP	Domain name resolution

Local users will use an internal domain name server, instead of having access

to an external DNS server.

Network Diagram

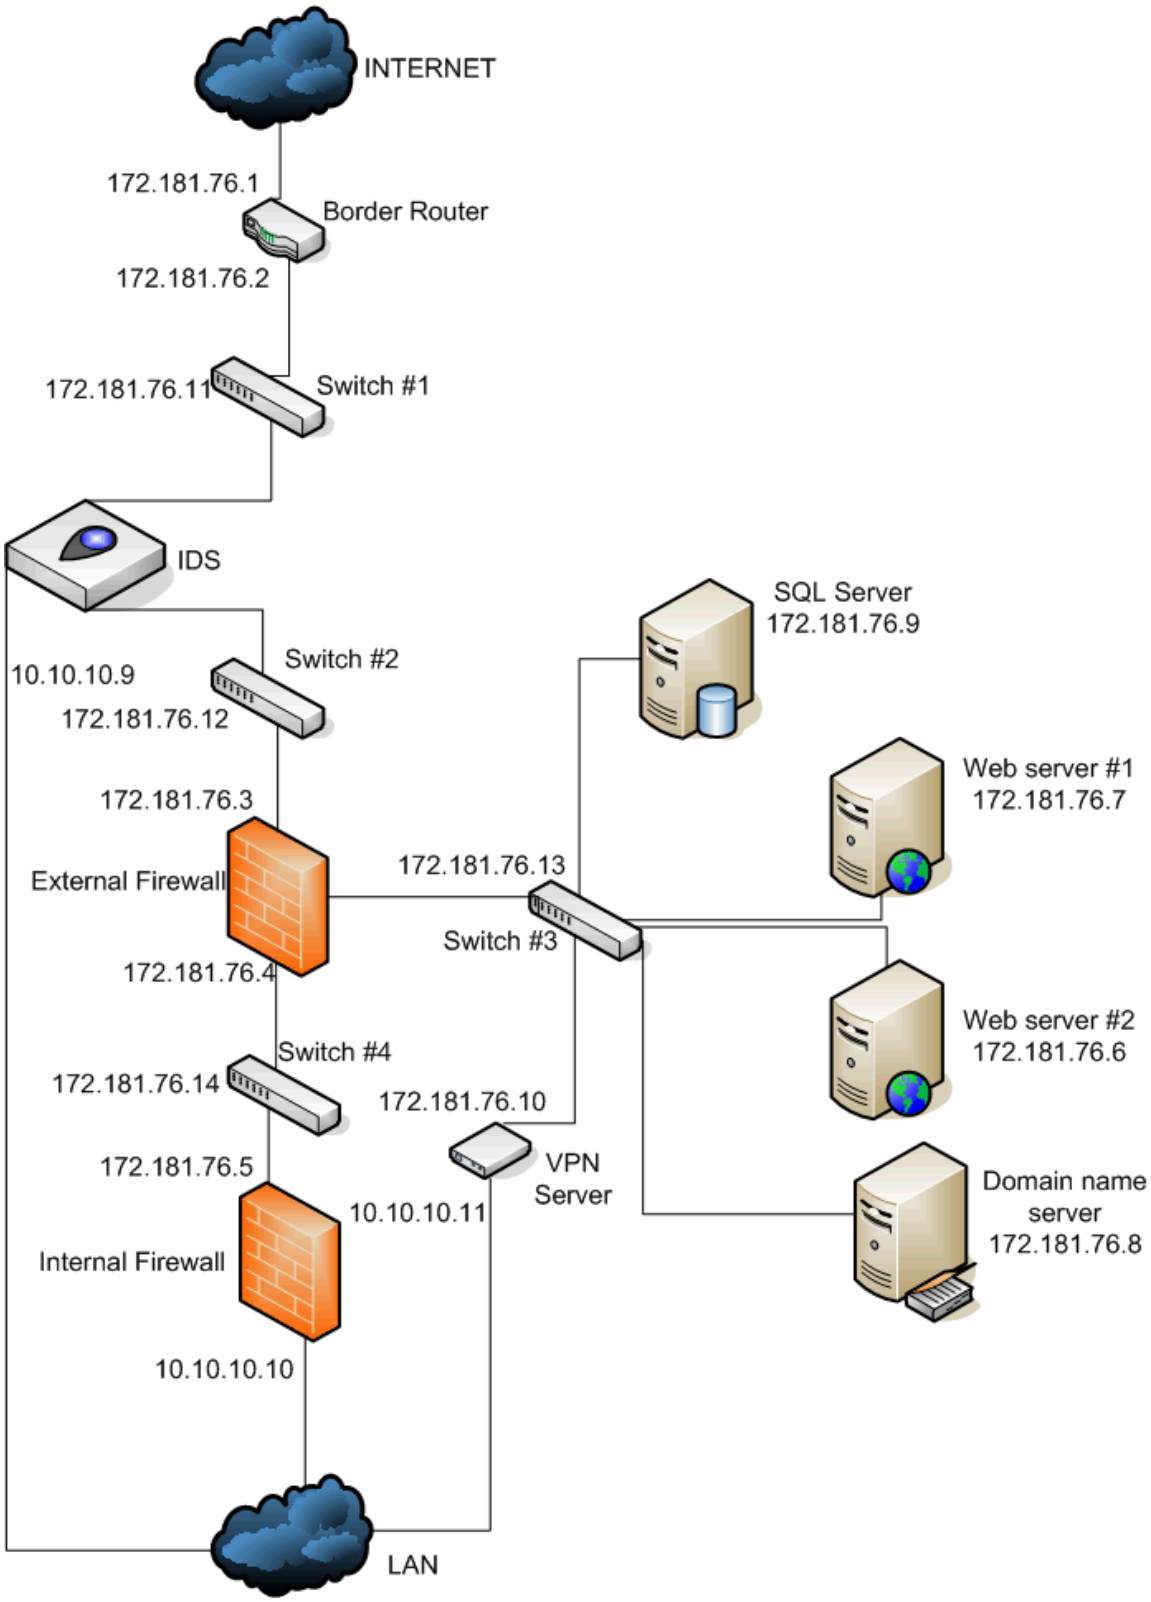

Defense-in-depth

The preceding network diagram illustrates how GIAC Enterprises uses defense-in-depth to secure its network. The basic definition of the defense-in-depth concept is to use multiple, and different layers, of security between a potential intruder and its target in order to reduce vulnerabilities and increase the difficulty of reaching the target. This concept is neither new nor particular to the IT industry since it was already used during the Middle Age by kings who built both a moat and a high wall around their castle.

As applied to network security, defense-in-depth generally means implementing different security devices between a (semi) public network and an information host. Not one device is fool-proof, nor is any expected to be, but their juxtaposition adds layers of complexity to any attempted unauthorized entry in the hope that the sum of the parts will be enough to deter an attack.

There is no set defense-in-depth model outlying the number and nature of devices to be used. The final architecture depends on the needs of the organization as well as the data path. For instance, if someone were to attempt to gain unauthorized access to one of GIAC Enterprises' web servers, he or she would have to traverse three security devices in order to do so, the border router, the IDS and the external firewall. If the would-be hacker's intentions were to penetrate the internal network, however, he or she would find the internal firewall in their path along with all preceding devices.

GIAC Enterprises' network could have been designed differently to add even more layers of defense by positioning a dedicated firewall between switch #4 and each of the hosts (database server, DNS server, web servers 1 & 2, VPN server) connected to this switch. The author chose not to include the additional firewalls in the network architecture however, because the expected pay-off isn't worth the additional management complexities five firewalls would add.

Security components

GIAC Enterprises' border router, intrusion detection system, firewall, VPN server and filtering router all have their specific function in the network's defense-in-depth design. While some devices may overlap in some of their functions, no two are exactly alike and they are combined to add a stronger level of security than any of them could provide individually.

Border Router

Product	Cisco 2610XM router with Cisco OS 12.1.2T and DS3/ATM module
Specifications	http://www.cisco.com/en/US/products/hw/routers/ps259/products_data_sheet0900aecd800fa5be.html
IP (external)	172.181.76.1
IP (internal)	172.181.76.2

As a border router, this device is both the first layer and the last layer of a defense-in-depth approach, depending on whether the traffic is inbound or outbound respectively. Routers today do more than just route traffic. Most routers, including the one above, now can participate in the security infrastructure of a network through Access Control Lists (ACLs). By examining traffic against its ACL, a router may drop a packet rather than send it off to its destination.

The exact configuration of this device is explained in section 3 of this document, but we can already mention that one of its critical purposes is to ensure that traffic that should not enter or leave GIAC Enterprises' network for any reason be dropped. Example of such traffic includes inbound packets originating from obviously spoofed IP addresses (such as addresses used on the internal network or addresses reserved for future use by the IANA).

Routers apply different sets of rules based on whether the traffic is inbound ("ingress filtering") or outbound ("egress filtering"), and it is therefore critical to insert the correct rule in the correct set. One weakness of most routers is that they are still stateless filters, which means every single packet traversing a router is checked against all rules until a match is found, irrespective of whether a particular packet is the first in a transmission or the reply to a query that has already transited the router. A practical consequence is that this router's egress filter may not be used to block outbound connections originating from one of the web servers on the network, because every single HTTP GET request would fail due to a lack of response.

In a perfect illustration of the defense-in-depth concept, placing a stateful packet filter (firewall) behind a router will address a router's lack of state-awareness because a firewall is fully aware of the state of a connection and can therefore block outbound connections originating from one of GIAC Enterprises' web servers without affecting HTTP GET answers.

This particular Cisco 2610XM router is part of the popular Cisco 2600 series of routers and is therefore well supported. It is also the smallest router series able to accommodate a T-3 connection while sporting enough memory to process egress and ingress filters without latency, making it an ideal candidate for GIAC Enterprises' needs.

Intrusion Detection System

Product	ISS Proventia G100 with SiteProtector
Specifications	http://www.iss.net/products_services/enterprise_protection/proventia/features.php
IP (console)	10.10.10.9

Technically speaking, the ISS Proventia G100 appliance is not an Intrusion Detection System. The security industry is slowly moving away from IDS to focus on IPS, or Intrusion Prevention Systems. Whereby an IDS would detect an intrusion and report it using one or more methods, an IPS will not only detect and report an intrusion, but will also act to prevent it. Attack prevention usually takes the form of dropping the offending packet, the entire connection or resetting it altogether (by sending a RST packet to the source).

IDS and IPS detect attacks by checking the actual payload for known exploits, suspicious non-compliance to RFCs or malicious fragmented packets. As such, they are able to block, among other things, injection attacks that take advantage of a weakness in a legitimate, public service on a host to feed it purposely malformed code and cause it to crash or, worse, allow its take over.

IDS and IPS complement firewalls and routers because of their ability to inspect the content of a packet, whereas the two other devices are limited to inspecting the state of a packet, or merely its origin and destination respectively.

While there are a number of locations where this IPS could have been located, its position between the border router and the firewall is the most effective one. Had it been placed in front of the border router, it could be overloaded with traffic that should never enter GIAC Enterprises' network to begin with and is better dropped by a router. The IPS could have been placed between the firewall and switch #3, but it would have been unable to inspect packets headed for the internal network (such has replies to HTTP GET queries). Likewise, if the IPS was moved to a spot between the firewall and switch #4, it would not be able to inspect any packet destined for GIAC Enterprises' public hosts and would leave them vulnerable to injection or worm-based attacks. Another option would have been to install two IPS', both behind the firewall and each unit would have a reduced load to handle since some traffic would be dropped by the firewall. This architecture would not influence the overall security environment of GIAC Enterprises, but would represent an additional cost that doesn't seem necessary, since the NSS Group test results indicate the Proventia G100 can process packets at wire speed.

The Proventia G100 is an appliance version of Internet Security System's reputed RealSecure Server Sensor software. In an independent test conducted by the NSS Group, the Proventia code base obtained excellent results, blocking

99% of all attacks out of the box, and an impressing 100% of attacks after a signature update, including so-called "false negatives" (known exploits that have been slightly modified to evade strict signature-based detection)[8].

The Proventia G100 operates as an inline bridge device and is therefore invisible to would-be hackers, making avoidance that much harder. Management is done through a console port directly connected to the internal network. While the Proventia G100's throughput isn't as high as that of the code base tested by the NSS Group, its 100 Mbps are sufficient to handle GIAC Enterprises' T-3 and the device may be, and should be, configured to drop excess traffic in case of network congestion. If an IDS is truly preferred to an IPS, the Proventia G100 can operate in a so-called inline simulation mode where it will behave as if it were an IDS, keeping all its features active save those used to prevent attacks.

ISS' SiteProtector software provides monitoring, reporting, alerting and remote management services for the Proventia G100.

External Firewall

Product	Cisco PIX 515E
Specifications	http://www.cisco.com/en/US/products/hw/vpndevc/ps2030/products_data_sheet09186a0080091b15.html
IP (external)	172.181.76.3
IP (internal)	172.181.76.4

A firewall is a stateful inspection filter, which means it examines each inbound and outbound packet within its context, rather than in a vacuum as a stateless router does. A firewall attains its stateful status by keeping track of active connections on the base of source and destination ports and IPs and protocol used. This allows a firewall to determine whether any given packet may be a reply to a certain query and treat it differently than if it were a query instead. For instance, a firewall may block all outbound connections from a given host, but packets that are sent by the host in reply to an inbound query will still be allowed through, whereas a stateless device doesn't track active connections and would block all outbound connections from the host, including those sent as a reply to a query.

The use of stateful inspection therefore makes a firewall a much more granular device than a router is, but the two complement one another as part of a defense-in-depth approach. As explained above, the router will be used to drop all inbound traffic that has absolutely no valid reason to reach GIAC Enterprises' network, thereby reducing the load on the firewall as it examines the state of every remaining packet. Because of its stateless nature, however, the border

[8] http://www.nss.co.uk/gigabitids/edition3/iss/iss.htm

router needs to grant every legitimate IP access to the LAN through its external IP address (172.181.76.5), irrespective of whether a packet bound for the local network is a query reply or not. The firewall's stateful inspection remedies to this obvious security threat by blocking all Internet traffic destined for the internal network, while relying on its context-awareness to still allow through query replies.

The purpose of this external firewall is to regulate traffic flows between the Internet, GIAC Enterprises' public hosts and the internal network. Exact rules for this firewall are listed in section 3, but the overall guiding principle of its configuration is to block everything that is not specifically allowed.

The major weakness of a stateful packet filter is that it does not inspect the payload of the packets it examines. If a firewall rules allows TCP/80 access to a host for web page serving, any packet with the right port and IP address as destination will be allowed through, even if its payload contains malicious code destined to crash or take over the web server. As explained above, intrusion detection or prevention systems (IDS and IPS respectively) remedy to this weakness by examining the content of the packet itself, and reporting on or dropping harmful packets before they can reach their target.

This particular Cisco router has three interfaces, one of which is designated as the DMZ interface and grants access to GIAC Enterprises' external hosts. Network Address Translation (NAT) could have been used to map external IPs to corresponding internal IPs on the DMZ link, but this is not necessary because GIAC Enterprises has enough routable IPs to accommodate all its hosts and these hosts are separated from the internal network.

VPN Server

Product Specifications	Cisco VPN 3020 Concentrator http://www.cisco.com/en/US/products/hw/vpndevc/ps2284/products_data_sheet09186a00801d3b56.html
IP (public)	172.181.76.10
IP (private)	10.10.10.11

While the two firewalls and even the border router could in theory play the role of VPN server, GIAC Enterprises will allocate this task to a dedicated device instead.

The Cisco VPN 3020 Concentrator has an encrypted throughput of 50 Mbps, enough to handle even the largest spikes in VPN traffic since the T-3 line connecting the network to the Internet has a maximum throughput of 45 Mbps. While this VPN server can handle both SSL and IPSec VPNs, only the latter will be enabled to give GIAC Enterprises' remote users the greatest access to the

internal network. IPSec tunnels will use Encapsulating Security Payload (ESP) for encryption, a choice that will have to be reflected in router and firewall rules.

All IPSec tunnels will be established on the public interface of the VPN server, and packets will be forwarded to the internal network over the private interface after decryption, bypassing both the external and internal firewalls. This represent a slight security risk as it leaves the internal network open to attacks coming from a compromised remote host connecting over VPN, but this risk may be mitigated by the use of a policy enforcement solution such as Sygate's Secure Enterprise[9] to ensure that all relevant patches are applied and anti-virus definitions kept up to date before granting connection to the internal network.

In order to strengthen security in the case of system theft, all VPN users are required to authenticate against an internal domain controller using RADIUS. RADIUS traffic will use the private interface of the VPN server and will therefore not need any specific rule in either firewall.

Internal firewall

Product Specifications	Cisco PIX 515E http://www.cisco.com/en/US/products/hw/vpndevc/ps2030/products_data_sheet09186a0080091b15.html
IP (external)	172.181.76.5
IP (internal)	10.10.10.10

Just like the external firewall, the internal firewall's guiding security policy principle is to block anything that is not expressly authorized. As a result, some rule overlapping between the two firewalls is unavoidable. In fact, both firewalls will have the same set of rules as apply to inbound traffic destined for the internal network's routable address. Where the two firewalls differ is in their rules for outbound traffic as well as the fact the internal firewall will have NAT and PAT enabled[10].

Network Address Translation (NAT) is used to give all outbound connections the same external IP rather than their unique internal, non-routable IP. This not only limits the number of routable IPs GIAC Enterprises needs to have, but also keeps the internal network structure shielded from the Internet.

Much like stateful packet filtering, NAT maintains a table of active connections where each internal IP address is mapped to a dynamically assigned dedicated

[9] http://www.sygate.com/products/sygate-secure-enterprise.htm

[10] Technically, PAT is a subset of NAT, and in this particular case, PAT will be used both for inbound and outbound connections. In order to avoid any confusion, the term NAT will be used to designate address translations applied to outbound connections (specifically single-address NAT, also known as PAT) and the term PAT will be used for inbound address translations.

port number. For every outbound packet, NAT records the source IP and port, then replaces them with the public IP and the port it assigned the connection so that the packet's origin seems to be the firewall's public IP to all other devices. When the firewall receives a reply, it extracts the destination port from the packet and looks up the corresponding internal IP address and port in its NAT table, inserts them in the packet and forwards it.

Port Address Translation (PAT) is enabled on this firewall for three specific services located on the internal network that need to be available from the outside. Each of the three services has its own dedicated port, so PAT maps an inbound connection on the firewall's external IP address using the service's specific port to the internal IP address where the service resides by replacing source and destination IPs in packets as described above for NAT. The three services GIAC Enterprises needs to make available on its internal network are its SMTP server (TCP/25), Syslog server (UDP/514) and SQL server (TCP/1433).

As the first filtering device that outbound connections encounter, the internal firewall's rules set for outbound connections will not be entirely similar to the rules applied by the external firewall. Much like the border router drops packets that should never reach GIAC Enterprises, the internal firewall can block connections that should never leave the internal network. Such connections include packets where the destination IP is obviously spoofed[11] or packets addressed to services that GIAC Enterprises employees are not authorized to access, such as instant messaging, real-time chat and audio streaming[12]. Finally, and just as importantly, firewall rules may also block access to ports used by well-known trojans[13].

Additional components

Each switch on GIAC Enterprises' network has its own IP address for management purposes. Firewall rules will have to be designed in order to preserve access to each switch's management interface. Some switches, for instance switch #1 and switch #2, could be replaced by a crossover cable, but the usage of a switch is preferred for diagnostic and monitoring purposes. Using port spanning, a probe may be connected to each switch in order to record traffic before and after the IPS as a mean to gauge its efficiency. Cisco Catalyst 2940 series switches will be enough to fulfill GIAC Enterprises' switching needs in the external network. These switches are manageable via telnet.

---

[11] Non-routable IP addresses and reserved IP addresses.

[12] Blocking some of these services is fairly easy because they use well-known ports (TCP/5190 for AIM), but other services use a number of ports that also have legitimate use (TCP/80 or TCP/21 for Yahoo Instant Messenger) and the only way to effectively block them using a firewall is to block access to all IPs where this service is available.

[13] http://www.doshelp.com/Ports/Trojan_Ports.htm

The two web servers, the SQL server and the DNS server all are running on Windows 2003 hardened as per Microsoft's recommendations for exposed servers[14]. These servers are configured as stand-alone servers without a domain in order to reduce traffic, and holes through the firewall. Remote management is performed via Remote Desktop Protocol. The SQL server will not be accessible directly from the Internet; rather the two web servers will initiate a connection to the SQL server to answer queries from customers, partners and suppliers alike. The SQL server will be updated by a SQL server on the internal network to which it will forward cookies orders and fortune submissions. In addition to being accessible to DNS clients on the Internet, the DNS server is also configured to allow zone transfers to two secondary DNS servers located at 64.12.51.132 and 64.12.187.24[15].

The firewall rules will have to reflect the various access requirements for these components.

[14] Microsoft, "Windows 2003 Server Security," Chapter 11: Hardening Bastion Hosts, April 2003 http://www.microsoft.com/technet/security/prodtech/windowsserver2003/w2003hg/sgch00.mspx
[15] In "real life," these two DNS servers are owned by AOL.

Section 3

In a well designed security policy, firewalls and routers complement one another as part of a defense-in-depth approach. A border router can be used to not only route the traffic between a corporate network and the Internet, but also regulate it. By using Access Control Lists (ACLs), traffic that is obviously spoofed can be dropped at the first point of contact with the corporate network, easing the load on other devices. Routers have their limitation, however, one of them being that they are generally still stateless devices, and must for instance accept all incoming traffic when only traffic that is a response to an outbound request should in fact be allowed in. Because the primary nature of a router is to route, routers also tend to fail open, which would present an unacceptable security risk if they presented the only filter between the Internet and an internal network.

Firewalls are stateful packet filters, able to make filtering decisions based on context information determined from the state of a connection. Stateful inspection requires more processing power than stateless inspection, however, given the need to memorize connection state, so the use of a router's ACL to limit access to the firewall will ease the load on the firewall. To be sure, certain filters should be included both on a router and on a firewall, not only due to the risk of a router failing open, but also because this presents an added layer of prevention against unauthorized penetration attempts. Conversely, other filters should be unique to one or the other. For instance, a firewall may need to allow outbound traffic from management consoles on the internal network to public servers, while the router would prevent traffic from the same consoles from reaching the Internet to preserve their integrity.

Border routers

Routers process ACLs in cardinal order, starting with the first one until a match is found. Once a rule is matched to incoming or outgoing traffic, the router processes the traffic according to the rule and does not read additional filters.

It is therefore critical that rules be listed in the appropriate order, because an erroneous ranking of filtering rules may lead to a wide-ranging Accept rule being triggered before a narrow Deny rule, while the opposite was intended.

Earlier routers used so-called Standard IP ACL, whereby only the source IP of the traffic was examined. Most routers in use today are at least able to process Extended IP ACLs, where destination IP, protocol, port (for TCP and UDP) and sequences (for ICMP) are also considered. The tables below assume the use of Extended IP ACLs.

Ingress filtering

The table below contains all ingress rules for GIAC's border router. They are processed in the order in which they are listed, starting with rule # 1.

#	A/D	Protocol	Source			Destination		
			Start IP	End IP	Port	Start	End	Port
1	D	Any	10.0.0.0	10.255.255.255	Any	0.0.0.0	255.255.255.255	Any
2	D	Any	172.16.0.0	172.31.255.255	Any	0.0.0.0	255.255.255.255	Any
3	D	Any	192.168.0.0	192.168.255.255	Any	0.0.0.0	255.255.255.255	Any
4	D	Any	172.181.76.1	172.181.76.14	Any	0.0.0.0	255.255.255.255	Any
5	D	Any	0.0.0.0	2.255.255.255	Any	0.0.0.0	255.255.255.255	Any
6	D	Any	5.0.0.0	5.255.255.255	Any	0.0.0.0	255.255.255.255	Any
7	D	Any	7.0.0.0	7.255.255.255	Any	0.0.0.0	255.255.255.255	Any
8	D	Any	23.0.0.0	23.255.255.255	Any	0.0.0.0	255.255.255.255	Any
9	D	Any	27.0.0.0	27.255.255.255	Any	0.0.0.0	255.255.255.255	Any
10	D	Any	31.0.0.0	31.255.255.255	Any	0.0.0.0	255.255.255.255	Any
11	D	Any	36.0.0.0	37.255.255.255	Any	0.0.0.0	255.255.255.255	Any
12	D	Any	39.0.0.0	39.255.255.255	Any	0.0.0.0	255.255.255.255	Any
13	D	Any	41.0.0.0	42.255.255.255	Any	0.0.0.0	255.255.255.255	Any
14	D	Any	73.0.0.0	79.255.255.255	Any	0.0.0.0	255.255.255.255	Any
15	D	Any	89.0.0.0	123.255.255.255	Any	0.0.0.0	255.255.255.255	Any
16	D	Any	127.0.0.0	127.255.255.255	Any	0.0.0.0	255.255.255.255	Any
17	D	Any	173.0.0.0	187.255.255.255	Any	0.0.0.0	255.255.255.255	Any
18	D	Any	189.0.0.0	190.255.255.255	Any	0.0.0.0	255.255.255.255	Any
19	D	Any	197.0.0.0	197.255.255.255	Any	0.0.0.0	255.255.255.255	Any
20	D	Any	223.0.0.0	223.255.255.255	Any	0.0.0.0	255.255.255.255	Any
21	D	Any	224.0.0.0	255.255.255.255	Any	0.0.0.0	255.255.255.255	Any
22	A	UDP	0.0.0.0	255.255.255.255	53	172.181.76.8	172.181.76.8	53
23	A	TCP	64.12.51.132	64.12.51.132	53	172.181.76.8	172.181.76.8	53
24	A	TCP	64.12.187.24	64.12.187.24	53	172.181.76.8	172.181.76.8	53

25	A	TCP	0.0.0.0	255.255.255.255	Any	172.181.76.7	172.181.76.7	80
26	A	TCP	0.0.0.0	255.255.255.255	Any	172.181.76.7	172.181.76.7	443
27	A	TCP	0.0.0.0	255.255.255.255	Any	172.181.76.6	172.181.76.6	80
28	A	TCP	0.0.0.0	255.255.255.255	Any	172.181.76.6	172.181.76.6	443
29	A	UDP	0.0.0.0	255.255.255.255	500	172.181.76.10	172.181.76.10	500
30	A	IP50 (ESP)	0.0.0.0	255.255.255.255	N/A	172.181.76.10	172.181.76.10	N/A
31	A	TCP	0.0.0.0	255.255.255.255	123	172.181.76.7	172.181.76.8	N/A
32	A	TCP	0.0.0.0	255.255.255.255	Any	172.181.76.5	172.181.76.5	Any
33	A	UDP	0.0.0.0	255.255.255.255	Any	172.181.76.5	172.181.76.5	Any
34	D	Any	0.0.0.0	255.255.255.255	Any	0.0.0.0	255.255.255.255	Any

Rules 1, 2, and 3:
These three rules together block out spoofed IP packets masquerading as coming from internal IP addresses as defined by RFC 1918[16]. Illicit packets often spoof internal IP addresses in an attempt to fool the destination host into considering the source as an internal host rather than an external one and exploiting privileges reserved for internal users.

Rule 4:
This rule blocks incoming traffic spoofing one of the IPs assigned to GIAC Enterprise. The purpose of spoofing one of those IPs is to masquerade illicit packets as originating from an authorized GIAC Enterprise external host when it is not.

Rules 5 through 15:
These rules block incoming traffic spoofing one of the IANA's reserved IP blocks. Spoofing an IP marked as reserved by the IANA is intended to cloak the true source of the traffic, but is less effective[17] than spoofing either an internal or external IP used by GIAC Enterprise.

Rule 16:
Rule 16 is similar in all aspects to rules 5 through 15, with the added benefit of blocking the loopback IP address. Spoofing the loopback IP address is intended into fooling the target host into considering the packets as originating from itself.

Rules 17 through 21:
The purpose of these rules is similar in all aspects to that of rules 5 through 15 but it also includes multicast IP addresses.

Rule 22:
Rule 22 is a network rule as its purpose is to allow traffic from DNS clients on the internet to query GIAC Enterprise's own DNS server (IP address: 172.181.76.8) in order to resolve host names within the GIAC domain. Domain resolution often generates a large amount of traffic since it is necessary to access other services available on the network by their host name, so this rule is configured to be processed before other rules granting access to external services.

[16] IETF Network Working Group, Address Allocation for Private Internets, February 1996, http://rfc.net/rfc1918.html
[17] From a hacker's point of view.

Rule 23:
Rule 23 allows DNS zone transfers from GIAC Enterprises' primary DNS server located at 172.181.76.8 to a secondary DNS server located at 64.12.51.132. While DNS clients use UDP to resolve a single host name at a time, DNS servers use TCP when transferring large amounts of data between themselves, such as copying the entire DNS table for the domain. This rule enforces GIAC Enterprises' stated policy of allowing DNS transfers only to its secondary DNS servers.

Rule 24:
Rule 24 is similar to the preceding rule, but it affects another secondary DNS server located at 64.12.187.24.

Rule 25:
Rule 25 allows access from any address on the Internet (except the ones blocked by rules 1 through 22) to GIAC's public web server (IP: 172.181.76.7) over HTTP.

Rule 26:
Rule 26 is similar to rule 25, but applies to HTTPS, or secure HTTP instead. Whereas HTTP would be used for most pages on the GIAC web site, the ordering process is protected by HTTPS, making this rule necessary.

Rule 27:
Rule 27's purpose is similar to that of rule 25, but grants access instead to the GIAC web server reserved for partners and suppliers.

Rule 28:
Rule 28 is equally similar to rule 26, but is also intended to allow partners and suppliers to be able to conduct secure transactions on the GIAC web server dedicated to them.

Rule 29:
This rule is intended for GIAC's remote workforce, be it mobile salespeople or telecommuters. It allows for the negotiations and deletions of Security Associations accomplished by ISAKMP in order to open or close IPSec VPN connections. Since mobile salespeople are, by definition, mobile, this rule will accept all source IPs as long as the destination IP is that of GIAC's VPN server (IP: 172.181.76.10). If GIAC had only telecommuters, and if all telecommuters were set up with a static IP, this rule could have been strengthened by the use of those static IPs as source IP, but there are few companies today who do not need to allow some of their staff to connect to the internal network via VPN from any given location.

Rule 30:
By allowing IPSec/ESP protocol traffic to the VPN server, this rule makes it possible for the above-mentioned remote users to establish a VPN tunnel from their client to GIAC's VPN server. If rule 29 didn't exist, this rule would be superfluous because the Security Associations necessary for IPSec/ESP VPN couldn't be negotiated.

Rule 31:
While not absolutely necessary, coordinating time on all hosts with a single, reliable source is desirable. This rule will allow NTP traffic to reach GIAC Enterprise's external hosts.

Rules 32 and 33:
The last rule, rule 34, blocks all sources of traffic not already blocked or authorized by the preceding rules. Since most routers today are still stateless, it is therefore necessary to specifically allow incoming traffic directed at the external IP address (172.181.176.5) of GIAC's internal network, otherwise this traffic would be blocked by rule 33 and GIAC's internal users would be unable to access any service across the Internet.

Rule 34:
As outlined above, the purpose of rule 34 is to block any incoming traffic that hasn't already been authorized or rejected. A filter such as rule 34 is essential for most border routers, but also makes the configuration of ACLs that much more critical. If it weren't for generic rejection rules such as this one, any traffic not already filtered would be allowed to continue unimpeded, an obvious security risk. The downside, however, is that if an acceptance rule is missing from the ACL, legitimate traffic will be rejected, which could mean that one of GIAC's external services will not be reachable, or in a worse case scenario, most services if the missing rule happens to be one similar to rule 23.

Egress filtering

#	A/D	Protocol	Source			Destination		
			Start IP	End IP	Port	Start	End	Port
1	D	Any	0.0.0.0	255.255.255.255	Any	10.0.0.0	10.255.255.255	Any
2	D	Any	0.0.0.0	255.255.255.255	Any	172.16.0.0	172.31.255.255	Any
3	D	Any	0.0.0.0	255.255.255.255	Any	192.168.0.0	192.168.255.255	Any
4	D	Any	0.0.0.0	255.255.255.255	Any	172.181.76.1	172.181.76.14	Any
5	D	Any	0.0.0.0	255.255.255.255	Any	0.0.0.0	2.255.255.255	Any
6	D	Any	0.0.0.0	255.255.255.255	Any	5.0.0.0	5.255.255.255	Any
7	D	Any	0.0.0.0	255.255.255.255	Any	7.0.0.0	7.255.255.255	Any
8	D	Any	0.0.0.0	255.255.255.255	Any	23.0.0.0	23.255.255.255	Any
9	D	Any	0.0.0.0	255.255.255.255	Any	27.0.0.0	27.255.255.255	Any
10	D	Any	0.0.0.0	255.255.255.255	Any	31.0.0.0	31.255.255.255	Any
11	D	Any	0.0.0.0	255.255.255.255	Any	36.0.0.0	37.255.255.255	Any
12	D	Any	0.0.0.0	255.255.255.255	Any	39.0.0.0	39.255.255.255	Any
13	D	Any	0.0.0.0	255.255.255.255	Any	41.0.0.0	42.255.255.255	Any
14	D	Any	0.0.0.0	255.255.255.255	Any	73.0.0.0	79.255.255.255	Any
15	D	Any	0.0.0.0	255.255.255.255	Any	89.0.0.0	123.255.255.255	Any
16	D	Any	0.0.0.0	255.255.255.255	Any	127.0.0.0	127.255.255.255	Any
17	D	Any	0.0.0.0	255.255.255.255	Any	173.0.0.0	187.255.255.255	Any
18	D	Any	0.0.0.0	255.255.255.255	Any	189.0.0.0	190.255.255.255	Any
19	D	Any	0.0.0.0	255.255.255.255	Any	197.0.0.0	197.255.255.255	Any
20	D	Any	0.0.0.0	255.255.255.255	Any	223.0.0.0	223.255.255.255	Any
21	D	Any	0.0.0.0	255.255.255.255	Any	224.0.0.0	255.255.255.255	Any
22	A	UDP	172.181.76.8	172.181.76.8	53	64.12.51.132	64.12.51.132	53
23	A	TCP	172.181.76.8	172.181.76.8	53	64.12.187.24	64.12.187.24	53
24	A	TCP	172.181.76.8	172.181.76.8	53	0.0.0.0	255.255.255.255	53
25	A	TCP	172.181.76.7	172.181.76.7	80	0.0.0.0	255.255.255.255	Any
26	A	TCP	172.181.76.7	172.181.76.7	443	0.0.0.0	255.255.255.255	Any

27	A	TCP	172.181.76.6	172.181.76.6	80	0.0.0.0	255.255.255.255	Any
28	A	TCP	172.181.76.6	172.181.76.6	443	0.0.0.0	255.255.255.255	Any
29	A	UDP	172.181.76.10	172.181.76.10	500	0.0.0.0	255.255.255.255	500
30	A	IP50 (ESP)	172.181.76.10	172.181.76.10	N/A	0.0.0.0	255.255.255.255	N/A
31	A	TCP	172.181.76.7	172.181.76.8	Any	0.0.0.0	255.255.255.255	123
32	A	TCP	172.181.76.5	172.181.76.5	Any	0.0.0.0	255.255.255.255	Any
33	A	UDP	172.181.76.5	172.181.76.5	Any	0.0.0.0	255.255.255.255	Any
34	D	Any	0.0.0.0	255.255.255.255	Any	0.0.0.0	255.255.255.255	Any

The egress filtering table is very similar to the ingress filtering rules, but with the source and destination largely inverted. The purpose is no longer to keep packets with spoofed source IPs out and to correctly filter authorized services to the appropriate servers, such as DNS, but it is to prevent traffic from leaving the corporate network bound for spoofed IPs, or to ensure that only authorized traffic will flow from external servers to the Internet.

The purpose of this first set of rules is to block traffic from the corporate network addressed to obviously spoofed IP addresses in a similar pattern as explained for ingress filtering.

If router ACLs generally allowed a range of ports per rule rather than a single port, rules 32 and 33 could be strengthened by allowing outbound traffic from the internal network from ephemeral ports only, since no server service on the internal network should be allowed to start an outbound connection.

External Firewall

While the entire border router ACL could be included in the firewall rules, some of them are instead left out in the interest of diminishing the load on the firewall. It is therefore up to the border router alone to block obviously spoofed IP addresses.

As explained in section 2, the firewall rules have to account for management and reporting traffic in addition to inbound connections to GIAC Enterprises' public hosts. This traffic takes the shape of telnet connections to the switches and Remote Desktop Protocol connections to the Windows Server, while reporting is being provided by Syslog back to the internal network.

#	A/D	Protocol	Source	Destination

			Start IP	End IP	Port	Start	End	Port
1	A	TCP	172.181.76.6	172.181.76.7	Any	172.181.76.9	172.181.76.9	1433
2	A	UDP	172.181.76.1	172.181.76.4	Any	172.181.76.5	172.181.76.5	514
3	A	UDP	172.181.76.11	172.181.76.14	Any	172.181.76.5	172.181.76.5	514
4	A	TCP	172.181.76.7	172.181.76.8	Any	0.0.0.0	255.255.255.255	123
5	A	TCP	172.181.76.9	172.181.76.9	Any	172.181.76.5	172.181.76.5	1433
6	A	TCP	172.181.76.8	172.181.76.8	Any	0.0.0.0	255.255.255.255	53
7	D	TCP	172.181.76.1	172.181.76.4	Any	0.0.0.0	255.255.255.255	Any
8	D	UDP	172.181.76.1	172.181.76.4	Any	0.0.0.0	255.255.255.255	Any
9	D	TCP	172.181.76.6	172.181.76.14	Any	0.0.0.0	255.255.255.255	Any
10	D	UDP	172.181.76.6	172.181.76.14	Any	0.0.0.0	255.255.255.255	Any
11	A	TCP	0.0.0.0	255.255.255.255	Any	172.181.76.5	172.181.76.5	25
12	A	UDP	0.0.0.0	255.255.255.255	Any	172.181.76.8	172.181.76.8	53
13	A	TCP	0.0.0.0	255.255.255.255	Any	172.181.76.7	172.181.76.7	80
14	A	TCP	0.0.0.0	255.255.255.255	Any	172.181.76.7	172.181.76.7	443
15	A	TCP	0.0.0.0	255.255.255.255	Any	172.181.76.6	172.181.76.6	80
16	A	TCP	0.0.0.0	255.255.255.255	Any	172.181.76.6	172.181.76.6	443
17	A	UDP	0.0.0.0	255.255.255.255	500	172.181.76.10	172.181.76.10	500
18	A	IP50 (ESP)	0.0.0.0	255.255.255.255	N/A	172.181.76.10	172.181.76.10	N/A
19	A	TCP	172.181.76.5	172.181.76.5	Any	172.181.76.1	172.181.76.4	23
20	A	TCP	172.181.76.5	172.181.76.5	Any	172.181.76.11	172.181.76.13	23
21	A	TCP	172.181.76.5	172.181.76.5	Any	172.181.76.6	172.181.76.9	3389
22	A	TCP	172.181.76.5	172.181.76.5	Any	172.181.76.9	172.181.76.9	1433
23	D	TCP	0.0.0.0	255.255.255.255	Any	172.181.76.1	172.181.76.14	Any
24	D	UDP	0.0.0.0	255.255.255.255	Any	172.181.76.1	172.181.76.14	Any
25	A	TCP	172.181.76.5	172.181.76.5	Any	0.0.0.0	255.255.255.255	Any
26	A	UDP	172.181.76.5	172.181.76.5	Any	0.0.0.0	255.255.255.255	Any
27	D	TCP	0.0.0.0	255.255.255.255	Any	0.0.0.0	255.255.255.255	Any
28	D	UDP	0.0.0.0	255.255.255.255	Any	0.0.0.0	255.255.255.255	Any

Rule 1:
This rule allows both web servers to initiate a connection with the MS SQL server on the internal network via the NAT address of the internal firewall. This rule assumes that the internal firewall has port forwarding enabled to redirect the connection to the SQL server's actual internal address.

Rules 2 & 3:
Rule 2 and 3 let firewalls and switches respectively forward their log entries to the syslog server located on the internal network by way of PAT running on the public IP of the internal firewall.

Rule 4:
This rule allows all external hosts to query a NTP server for the exact time.

Rule 5:
Rule 5 allows the external SQL Server to initiate a connection with the SQL Server on the internal network through PAT on the internal firewall's public interface.

Rule 6:
Rule 6 allows the primary DNS server to push zone updates to secondary DNS servers. While this rule allows traffic to flow freely to any IP address, a rule in the router's egress filtering table ensures that only authorized secondary DNS servers receive domain updates.

Rules 7 through 10:
These four rules combined prevent all external hosts, except the public IP address of the internal firewall, from initiating any connection to any address, other than the ones allowed by rules 1 through 6 already. Since a firewall is state-aware, these two rules will not impede replies to queries received from other hosts, be they on the corporate network or on the Internet.

Rule 11:
Rule 11 allows SMTP traffic to reach the external IP of the LAN where it will be redirected to the appropriate internal server using PAT.

Rules 12 through 18:
These rules have a similar objective to their counterparts in the border router's ingress filters, namely to allow clients inside and outside the corporate network to connect to GIAC's external hosts for specific purposes (name resolutions, HTTP and HTTPS browsing and VPN connections).

Rule 19:
This rule allows the use of telnet to manage the border router and the external firewall from the internal network.

Rule 20:
Rule 20 is similar to rule 19, but in regard to switches instead.

Rule 21:
Rule 21 makes it possible to use RDP to manage the four Windows servers from the internal network.

Route 22:
Route 22 allows the SQL Server on the internal network to initiate a connection with the external SQL Server.

Rules 23 & 24:
These two rules block all traffic to the external hosts, unless it has already been authorized by a preceding rule.

Rules 25 & 26:
Since the last two rules prevent all traffic from transiting through the firewall, rules 25 and 26 are necessary to allow outbound traffic from the internal network. Arguably, these rules could be strengthened to ensure that outbound traffic is for business purposes only by one of several ways.
Firewall rules could be re-written to grant access only to ports that are necessary for business purposes (TCP/80 for web browsing for instance), or to deny access to services that are definitively not business related (AIM at TCP/5190 for example) but these rules are better applied on the internal firewall, since it is the first filter that outbound traffic encounters.

Rules 27 and 28:
As explained above, these two rules block any traffic that hasn't been specifically allowed by one of the preceding rules. The purpose of these final two rules is to serve as a catch-all for a maximum security policy. The alternative would be not to include any rule, with the result that any traffic that hadn't been blocked would be allowed, since firewalls allow all traffic by default[18]..

Conclusions

As with border routers, firewalls process rules in a cardinal order and stop when reaching the first matching one, making the rule ranking just as critical. Generally speaking, Allow and Deny rules affecting only a specific host, or a small number of hosts, should come before rules affecting an entire subnet, otherwise they may not be accessible at all.

This is why rules 2 through 4 could not have been ranked anywhere else in the table above. The intent was to prevent external hosts from initiating any type of

[18] Most firewalls come pre-configured with a Block-All rule anyhow.

traffic except for the purpose of transferring syslog data, and adjusting their clock to a common time server. In order to achieve this objective, four blanket rules (7 – 10) blocking outbound TCP and UPD are necessary, but they had to be preceded by rules allowing connections originating from the external hosts, and they had to come before any rule allowing all incoming traffic to an external host, even if for specific ports. Without this careful ranking, external hosts might not have been able to adjust their internal clock, or they may have been prevented from generating outbound traffic destined for the Internet, but not for one another.

This example illustrates the critical role that ranking plays in the effectiveness of ACLs, both in firewalls and in routers. To be sure, ranking isn't always this critical. It doesn't matter, for instance, whether firewall rule 3 precedes or follows firewall rule 4, because neither affects the other, since one filters TCP traffic and the other UDP traffic.

A similar example can be found with router ingress rules 1 through 21. Spoofed IP traffic should be dropped as soon as possible to ensure that it isn't included in an Allow rule by accident, but whether or not rule 1 precedes rule 2 is immaterial as long as both rules come before any rule not related to spoofed IP.

The conclusion one may draw is that, for security purposes, the ranking of each rule isn't as critical as is the ranking of the group in which it may belong. We can usually place all rules in one of three groups.

One group includes all rules dropping spoofed IP packets as well as any other IP address that is unwanted on the corporate network for any reason. For example, many companies who have no business ties to certain Asian countries where most of the spam now comes from have included the IP ranges assigned to said countries in this group, ensuring that spam never even reach their mail server. Since IP addresses targeted by the rules in this group are unwanted under any condition, this group is generally first among router and firewall rules.

A second group unites together so-called network rules, or infrastructure rules. These rules generally affect traffic not directly related to user activity, but in support of it, such as DNS queries or database look-ups performed by a web server.

The third group covers user rules, that is rules filtering traffic directly generated by users, such as requesting a web page.

Splitting all rules among one of the three groups in the planning stage is strongly encouraged in order to ensure all necessary rules are included and avoid confusions. Rules may also be ranked in an ACL by group order, a method that has the advantage of making rule management easier.

Another method, preferred by the author, consists of putting all group 1 rules first, then ranking group 2 and 3 rules together according to their frequency of usage as well as the final security objectives. The router ingress filtering rules are a good example of this choice. If the rules were ranked strictly by group, rules 25 through 28 would be ranked below rules 29 through 31. It is the author's opinion that the current order is preferable, however, because it ensures that spoofed IPs are dropped immediately, then the remaining rules are examined against incoming traffic according to their likelihood of being triggered, saving processing time all the while maintaining network security. This being said, because of the lack of group ranking, managing such a filter set becomes exponentially complex as the number of rules increases. It is thus up to each network administrator to decide which methods suits him and the circumstances best.

References

Scott Foster, “Utilizing Static Packet Filters to Enhance Network Security,” October 22, 2004.
http://www.sans.org/rr/whitepapers/firewalls/1518.php

Eric Peeters, “Wireless Security Beyond WEP And WPA,” May 2, 2004.
http://www.sans.org/rr/whitepapers/wireless/1425.php

Tunix, “Firewall White Paper,” 2004.
http://www.tunix.nl/fire/wp-eng/

IANA, Internet Protocol v4 Address Space, January 2005.
http://www.iana.org/assignments/ipv4-address-space

IANA, Port Numbers, February 2005.
http://www.iana.org/assignments/port-numbers

Paul Funk, “802.11i secures Wireless LANs,” Network World, March 28, 2005.
http://www.nwfusion.com/news/tech/2005/032805techupdate.html

Diana Kelley & Lisa Phifer. “802.11 Planet - WLAN security tutorial,” June 2003.

Kevin Fogarty, “Eyes on 802.11i,” Network World, March 14, 2005.

Trapeze Networks, “Detecting Rogue Users and APs In A Wireless LAN.”
http://www.trapezenetworks.com/technology/whitepapers/detectingrogue.asp

Louise McKeag, “Access Points – Beyond The Thin/Fat Spat,” TechWorld, April 8, 2004.
http://www.techworld.com/mobility/features/index.cfm?FeatureID=477