Global Information Assurance Certification Paper

Deploying Honeypots and the Security Architecture of a Fictitious Company

GIAC Certified Firewall Analyst (GCFW)

Practical Assignment (v4.1)

David Pérez Conde

February 2005

ABSTRACT

This paper constitutes the practical assignment (v4.1) that I submitted as one of the requirements to obtain the GCFW certification (GIAC Certified Firewall Analyst).

It is divided in three parts: first, a short article discussing the benefits and implications of deploying honeypots in the network is included, second, the security architecture of a fictitious company is described, and finally, the firewall policy of the primary router and firewall in that architecture is presented in detail.

This page intentionally left blank

Table of Contents

1 Deploying Honeypots, Honeynets or Tar Pits

1.1 Abstract

This paper discusses the possibility of integrating honeypots in the overall security architecture of any organization.

After defining the terms *honeypot*, *honeynet* and *tar pit,* different kinds of honeypots are identified and classified: high-interaction vs. low-interaction honeypots. To better illustrate the concepts being discussed, a couple of real-life examples of honeypots, namely "honeyd" and "The Honeywall CD-ROM", are described.

Then, the benefits and implications (technical, ethical and legal) of honeypots are explored and finally a brief outlook at future applications, namely honeypot-based intrusion detection systems (HPIDS) and the use of honeytokens.

Finally, the reader is offered some conclusions.

1.2 Introduction

Honeypots are a relatively new security technology that has unique benefits to offer. However, they also present potentially unique implications that should be fully explored and understood in order to deploy honeypot based solutions in a proper manner.

1.3 Definitions

Before *honeypots, honeynets* and *tar pits* can be discussed, those terms must be defined.

A **honeypot** is defined by Lance Spitzner[1] [SPZ01] as "an information system resource whose value lies in unauthorized or illicit use of that resource". This is a broad definition, that allows for very different types of honeypots to be included in it.

Perhaps the most commonly known type of honeypot is "a computer connected to the network with the only purpose of being probed and attacked". But the definition allows for other kinds of honeypots to be included in the term, like the so-called *honeytokens,* small pieces of information planted on information systems with the only purpose of being illicitly accessed or modified [SPZ03].

A **honeynet** is a special type of honeypot: a specific network segment populated with, and only

1 Actually, this definition was the final result of long discussions that were held in the Honeypots mailing list[SPZ02] with the objective of finding the best definition possible for the term "honeypot".

with, honeypots [HON01]. Since a honeynet only contains honeypots, any interaction with it, that is any traffic going into or out from the honeynet, will most probably[2] correspond to malicious or unauthorized activities.

Finally, a **tar pit**[3] is a special type of honeypot that when attacked it reacts by taking up resources from the attacker system so as to slow it down or even halt it to a crawl.

1.4 Classification

Honeypots can be classified, according to the level of interaction of the attacker they allow, into low-interaction and high-interaction honeypots [HON01]. This classification is important because honeypots with a different level of interaction provide different quantity and quality of information and thus are suitable to serve different purposes. The benefits and implications of each type will be analyzed later.

A piece of software capable of emulating network services would be an example of low interaction honeypot. Attackers connecting to it would get responses similar to the real services but they would never be able to abuse these fake services using a exploit they may have coded for their real counterparts.

A computer system with a complete operating system (OS) installation and real network services running on it would be an example of high interaction honeypot. In this case, attackers connecting to it would get their responses from the real services and they might be able to break into the system and take over it should these services present any vulnerability.

Honeynets are the ultimate high-interaction honeypots: a complete network segment full of honeypots[4] ready for attackers to interact with.

1.5 Real-Life Examples

Several honeypot-based solutions are available in the market today. Some of them are commercial, like KFSensor [KFS01], Symantec Decoy Server [SYM01], or Specter [NET01], while others are free[5] software, like Honeyd [PRO01] and The Honeywall CDROM [HON03].

The three commercial examples given and Honeyd fall in the category of low-interaction honeypots: they all simulate systems and services. On the other hand, The Honeywall CDROM is a tool that allows the installation and configuration of honeynets in a very simple and effective way.

2 Sometimes the traffic may be simply generated by misconfigured systems.

3 The first famous tar pit in this honeypot sense was LaBrea, by Tom Liston [LIS01]. Its functionality has later been incorporated in honeyd [PRO01].

4 Usually high-interaction honeypots, although they could also be low-interaction or a mixture of both types.

5 Honeyd is distributed under the GPL license (http://www.gnu.org/copyleft/gpl.html) and The Honeywall CDROM contains software under the GPL and the BSD (http://www.opensource.org/licenses/bsd-license.php) licenses.

The latest two will be described in more detail in order to illustrate with real-life examples what low- and high- interaction honeypots may look like.

1.5.1 Honeyd

In the words of its creator, Niels Provos[6], "Honeyd is a small daemon that creates virtual hosts on a network. The hosts can be configured to run arbitrary services, and their personality can be adapted so that they appear to be running certain operating systems. Honeyd enables a single host to claim multiple addresses on a LAN for network simulation."

The number and type of virtual hosts for Honeyd to emulate is defined by the user in a configuration file. Honeyd is able to simulate not only some services running on the virtual systems but also the whole TCP/IP stack of those systems, so that they respond to OS fingerprinting (using nmap [FYO01], xprobe[XPR01], p0f [ZAL01], or similar tools) in the same way as the emulated OS would.

The feature to claim multiple IP addresses on a single host is very useful. A network range can be specified in the configuration file and Honeyd will assume any unused IP addresses in that range. This way, it will appear to an attacker scanning a whole class C network, for example, that all 254 IP addresses are used by that many hosts.

Another neat feature integrated in Honeyd is the possibility of behaving as a tar pit[7]. The idea, to the best of my knowledge, was first implemented in the LaBrea [LIS01] tar pit, by Tom Liston. The stickiness of the tar pit is achieved using a TCP option: the window size. The tar pit will accept any incoming connections, but as soon as the 3-way handshake is completed it will send a packet to the other end of the communication with a window size of zero. This will be interpreted by the other host (the attacker's) as the following message: "I really want to talk to you but I'm so busy right now!, please hold the line (do not drop the connection) and I will be back to you as soon as I can"[8]. The attacker system will then hold the connection, wait for a few seconds and send a single packet to the tar pit asking if the conversation can continue. The tar pit will respond the same as before and this procedure will repeat itself endlessly[LIS01]. This will keep some amount of memory occupied in the attacker system and will consume a few CPU cycles every now and then.

The tar pit feature, together with the ability to simulate systems for whole ranges of unused IP addresses, gives honeyd the ability to potentially slow down the spread of worms.

1.5.2 The Honeywall CDROM

In the words of their developers, "the Honeywall CDROM combines all the tools and

6 http://www.honeyd.org

7 The tar pit feature was introduced in version 0.7 of Honeyd.

8 Expression borrowed from Mike Poor, explaining the tar pit concept while teaching Track 3 (SECURITY 503: Intrusion Detection In-Depth) at SANS Computer Security 2003 (Monterey, CA ~ June 11-16).

requirements of a GenII honeynet gateway on a (hopefully) easy to use, secure, bootable CDROM. The intent is to make honeynets easier to deploy and customize. You simply boot off the CDROM, configure it based on your environment, and you should have a Honeywall gateway ready to go."

The term[9] "GenII honeynet gateway" (a.k.a. Honeywall) refers to the system that should always sit between any second generation honeynet and the rest of the world, being in charge of data capture (logging every packet entering or leaving the honeynet), data collection (receiving and storing information about the attackers activities coming from modified software on the honeypots), and data control (ensuring that attackers can not launch successful attacks from any compromised honeypot to the outside world). A second generation honeynet differs from a first generation honeynet mainly in the better data control features and higher level of stealthiness of the gateway. For example, a GenI honeywall would be a level three device (a firewall) limiting the number of outgoing connections that could be initiated from the honeynet, while a GenII honeywall would be a level two device (a bridge) with some kind of intrusion prevention system that would allow outgoing connections but would disable any attacks launched to the outer world by modifying the packets' payload on-the-fly. A level three device can be detected because of the decrement on the TTL (time to live) of packets it must perform. However, a level two device does not need to modify any passing packets at all, making a long way towards invisibility.

The Honeywall CDROM makes it very easy to deploy a honeynet. A system booted from this CDROM is almost automatically[10] configured as a full-fledged Honeywall gateway. The only thing left to have a fully functional honeynet working is to connect a set of honeypots in the "inside" of the Honeywall.

1.6 Benefits

Honeypots, either low- or high-interaction, can be an extremely powerful technology to be integrated in any overall security architecture. In particular, they are specially well suited to detect and record sources and types of known and unknown probes and attacks.

Anyone having worked with network-based intrusion detection systems (NIDS), which are supposed to fulfill the mission of alerting on network attacks, knows that they face two main problems: false positives, alarms triggered by unimportant events mistaken as attacks, and false negatives, real attacks not being reported. Honeypots, on the other hand, excel on these two areas. For one thing, because honeypots serve no real production purpose other than being attacked, any interaction whatsoever with them is by definition illegitimate traffic that should be

9 For more information on these terms and on honeynets in general, the reader is strongly encouraged to read the book "Know Your Enemy: Learning about Security Threats, (2nd Edition)" [HON02], by The Honeynet Project.

10 The few configuration options that need to be set, can be configured using a simple text-based menu that is executed the first time the Honeywall is booted.

reported and analyzed, leaving very little room for false positives[11]. Also, because the definition of attacks to alert on corresponds to any traffic entering or leaving the honeypot as opposed to be based on known attack signatures, honeypots would report even on brand new attacks no one has heard of before. Yet another cause of false negatives on NIDS systems sometimes is the very high network load these systems must cope with, which poses high requirements on resources (mainly CPU power) for NIDS systems not to start dropping packets. Honeypots, on the other hand, only have to deal with rogue traffic flowing to or from them or non-existent systems [PRO01], so even a relatively modest honeypot will usually be far from being overloaded even in a high load network.

This is not to say that honeypots should replace NIDS systems: on the contrary, they complement each other. Both technologies are able to detect attacks, but neither of them is able to detect all attacks. Some limitations of NIDS have been exposed. Honeypots also have theirs. Notably, they only detect attacks directed to themselves or to non-existent systems, being completely unable to alert on attacks directed to production systems.

The amount and type of information about the attacks detected by honeypots will vary depending on the type of honeypot.

A low-interaction honeypot will be able to provide the full contents of the packets received from the attackers, including their source IP address, the port or ports they are targeting, and the payload of these packets which may contain any exploits used in the attack only if the fake services were enough to fool the attacker into launching the full attack. Thus, the information that low-interaction honeypots can provide can be very useful but it is somewhat limited.

High-interaction honeypots, on the other hand, can provide much more information. Since they are running real services, it is much more probable that attackers launch their full attack and this will be recorded by the honeypot. Furthermore, if the attack is successful and the attackers break into the honeypot, they will probably modify it to keep access beyond the vulnerability they exploited to get in, try to cover their tracks, and then engage in whatever activities they might want the conquered system for (launching new attacks, setting up IRC servers, storing warez, etc.). All these activities will be recorded and reported by a properly configured high-interaction honeypot.

Honeynets expand the information provided by a single high-interaction honeypot to that of a set of different honeypots probably configured with different operating systems, different applications, and different versions of software, thus increasing the amount of information that can be gathered from different network attacks.

Tar pits are a special functionality that can be present in either low- or high- interaction honeypots and they can provide extra benefits. Although it has not yet been tested in production

11 Alerts on traffic generated by misconfigured systems may be considered false positives. However, they can also be considered true positives because although they don't reveal malicious activity they are in fact events of interest: they may well be the only way to detect that misconfiguration.

networks, theory and some lab tests have it that tar pits could help to slow down or even stop the propagation of worms that use TCP[12] (Transmission Control Protocol) to propagate [SPZ01]. When a worm is spreading, each infected system tries to infect as many other systems as it can, as fast as it can. If such an infected system tries to infect a tar pit, this connection will get stuck and some resources of the attacking system (mainly memory) will become locked in use. If the infected system happens to establish not only a single connection to a single tar pit but many connections to many tar pits[13], it may reach a point in which all of its resources will become locked and it will be unable to infect any more systems.

1.7 Implications

Deploying honeypots, however simple it may seem at first sight, must not be done without fully evaluating its implications. There are technical, ethical and legal aspects that should be taken into account while considering the deployment of honeypots. These aspects are discussed in the following sections.

1.7.1 Technical Implications

One of the most important factors that must not be overlooked is the manpower needed not only to configure and install the honeypots but also to maintain it and most importantly to analyze and act upon its alerts. A honeypot will be of little or no use if once installed it is left completely unattended because the security personnel does not have the time to look at the information it provides. This is not a new problem. Many NIDS deployment projects have failed miserably not because the NIDS product didn't work or because the configuration wasn't right[14] but because after the deployment there was nobody with the time required to go through the alerts and act on them.

Another important decision that must be made early in the design phase is where to locate the honeypots in the network, which will determine the kind of attacks they will detect. Honeypots in the DMZ, for example, being exposed to external traffic will detect external attacks and probes. Given the current amount of noise in the Internet this will probably amount for lots of unimportant probes and scans together with the important ones. On the other hand, honeypots in the internal network would detect internal attacks, either true malicious activity or just bad traffic generated by infected or misconfigured systems. The signal-to-noise ratio[15] in this case would probably be much higher. Likewise, a honeypot located in a servers' network will pick up different kinds and amounts of attacks than a honeypot in a users' network.

Last, but not least, the type of honeypot to be deployed must be considered carefully. A low-interaction honeypot will be easier to maintain but the information it will provide about each attack

12 Transmission Control Protocol [STE01]

13 It will be shown when “honeyd” is presented in the next section that this will happen in many cases.

14 Many projects have also failed because of this exact reason.

15 Number of important events divided by the number of unimportant events.

will be very limited. A high-interaction honeypot will be harder to configure and maintain but the information about the attacks will be much more detailed. Also, a high-interaction honeypot will introduce a certain level of risk: if the honeypot itself is compromised it could be used to launch new attacks from there. This extra risk must be considered and countermeasures must be designed to manage it.

All these factors must be taken into account when designing a honeypot deployment and the advantages and disadvantages of each option must be balanced against the objectives of the project.

1.7.2 Ethical Implications

There seems to be a general belief that deploying honeypots may have lots of legal and ethical implications. While legal implications certainly exist and must be addressed, as discussed in the next section, it is my humble opinion that there are no tough ethical implications regarding the deployment of honeypots.

The main concern could be the potential invasion of privacy of individuals communicating with or through a honeypot because potentially all the communication will be monitored and most probably recorded. However, by definition of a honeypot, any individual conducting such communication would be an unauthorized user of the honeypot. Now, does a rogue user have the right to have his or her communication through the honeypot considered and respected as private? Maybe legally he or she has, which I doubt, but ethically speaking I think he or she waived that right the very moment he or she accessed the honeypot without authorization.

The only special case that I think deserves some thought is: what if an attacker takes over a honeypot and uses it to store or circulate third party confidential information or evidence of a crime? My view on this is that the problem would not lie in the fact that the information was captured by the honeypot and therefore is made available to the honeypot administrator, but in the way the administrator deals with the information once he or she discovers its nature. In this regard, I don't think the use of a honeypot is any different than when we receive by mistake an e-mail message that wasn't intended to us. It is the way we react to that situation that will be either ethical or unethical. The possibility of this situation should be considered beforehand and the desired reaction written on policy.

1.7.3 Legal Implications

This topic is discussed at great level of detail by Richard Salgado, former Senior Counsel in the Computer Crime and Intellectual Property Section of the United States Department of Justice, on chapter 8 “Legal Issues” in the book “Know Your Enemy, 2nd Ed.” by The Honeynet Project [HON02]. This specific chapter is freely available online at the web page of the Honeynet Project: “http://www.honeynet.org/book/Chp8.pdf”.

In summary, finding out which laws apply to a particular honeynet deployment and to which extent is critical in order to avoid possible legal risks. For one thing, monitoring the network traffic of the honeypots may step over some legal rights of the users even if they are all unauthorized users. Second, if the honeypot is used by an attacker to commit a crime, the information gathered from that honeypot may require a very special treatment. And last, but not least, if a honeypot is used by an attacker to cause damage to third parties then those third parties may decide to litigate against the owner of the honeypot for economical compensation. Knowing exactly which laws would apply to those situations is a matter to be taken seriously and getting proper legal counsel before deploying a honeypot is highly recommended[16].

1.8 Future Applications

Honeypots are a relatively new technology and as such their full potential is yet to be explored. So far their main application has been the research of the methods and tactics of attackers, but honeypots have much more to offer.

Using honeypots as the base for a new kind of intrusion detection systems, "honeypot-based IDS (HPIDS)" [PER01], is one of the applications that should get the honeypots into the overall security architecture of production networks. Honeypot-based IDS (HPIDS) should be integrated with their system (HIDS) and network (NIDS) counterparts to conform a much more powerful intrusion detection architecture.

The use of honeytokens is another area that will probably see a boost in the near future, specially where intellectual property rights are the biggest concern. Planting bogus information on a database, for example, may serve as proof of unauthorized access to the database if that information is later found somewhere else.

1.9 Conclusion

Honeypots offer benefits that no other technology can provide. Although honeypots do not directly protect any systems, they improve the overall security posture of a network by allowing the network administrators to quickly identify systems from which attacks are being launched (for example, systems infected with a virus or worm) and providing them with information about the techniques used on those attacks. The network and system administrators can then apply that knowledge to identify production systems that may have been compromised using the same kind of attack and to recover from such incidents.

Different types of honeypots have different implications and benefits and they all should be considered together on each situation in order to choose the right honeypot solution for a given purpose and a given set of circumstances.

16 Actually, the same holds true for traditional intrusion detection systems (IDS) systems since they may incur the same legal risks. Or, the other way around, honeypots are not much more dangerous than IDS, legally speaking.

2 Security Architecture of "GIAC Enterprises"

2.1 Introduction

GIAC Enterprises is a fictitious company that sells fortune cookie sayings through the web. In this section the access requirements and restrictions of several groups of people interacting with GIAC are analyzed, the security architecture designed to meet those requirements and keep the shop as secure as possible is described and finally the filtering rules for the main router and firewall are detailed.

2.2 Access Requirements and Restrictions

2.2.1 Customers

Customers will purchase bulk online fortunes through a dedicated web server which will be called "Web Server for Customers" or "WCust" for short. Having a web server dedicated for customers reduces the chances of the web server being compromised by attackers because access can be restricted to a very reduced set of services. It also reduces the risk of the customers being impacted in the event of some other server being compromised (e.g. the web server exposed to the general public).

Customers must be able to authenticate the web server (WCust) to make sure they don't send their orders and payment data to a bogus web server planted by a third party, and vice versa, the web server must be able to properly authenticate each customer in order to give them access to their own data and only to their own data. Customers must be prevented from accessing other customers' data. The authentication of the server will be achieved by using SSL and a certificate signed by a well known certification authority. Authentication of the customers will be achieved via username and password which will be checked by the server against a RADIUS server. The access of each customer to their own data and only to that data will be guaranteed by the web application.

The fortunes must travel encrypted through the Internet to avoid the risk of being intercepted by third parties that could then resell them at a lower price. This will be achieved using SSL.

Through the web server (WCust), customers will be able to place orders, download fortunes and check the history of their past orders. The web application on the web server will allow all this functionality once the user has been properly identified.

Source	Destination	Port(s)/Protocol	Description
Customers	WCust	443/TCP (SSL)	Customer access to their data in order to place orders, download fortunes and check past orders.

Table 1 Access requirements for customers

2.2.2 Suppliers

Suppliers will upload fortune cookie sayings through a dedicated web server, which will be called "Web Server for Suppliers" or "WSupp" for short. The use of a dedicated web server for suppliers is recommended for the same reasons as for customers. On top of that, it helps in reducing the risk of suppliers and customers finding out about each other. If GIAC suppliers got to contact GIAC customers and sell them fortunes directly the whole GIAC Ent. business would be ruined.

The same access requirements and restrictions explained above for customers apply to suppliers.

Additionally, GIAC must be able to place orders on the suppliers. This will be achieved by allowing WSupp to initiate outgoing connections to and only to some specific suppliers' web servers.

Source	Destination	Port(s)/Protocol	Description
Suppliers	WSupp	443/TCP (SSL)	Supplier access to their data in order to upload fortunes and check past orders.
WSupp	Specific Suppliers' web servers	443/TCP (SSL)	GIAC access to suppliers' web servers in order to place orders.

Table 2 Access requirements for suppliers

2.2.3 Partners

Partners are international companies that translate and resell fortunes. With regards to GIAC, partners are just a special kind of customers, yet they should access their own dedicated server ("Web Server for Partners" or "WPart") in order to avoid interference with customers and the general public.

Other than that, partners will have the same access requirements as customers: only the web application will handle different information from them like their assigned discount level.

Source	Destination	Port(s)/Protocol	Description
Partners	WPart	443/TCP (SSL)	Partners access to their data in order to place orders, download fortunes and check past orders.

Table 3 Access requirements for partners

2.2.4 Employees on the internal network

Employees connected to the headquarters' internal network ("HQ Internal") need access to the internal servers (incoming e-mail, file, and application servers) POP3S, IMAPS, CIFS and a proprietary protocol.

They also need some Internet access: they need to browse the web (HTTP and HTTPS), download files using FTP and send e-mail (SMTP) to the Internet. This access is granted by a web and FTP proxy server, an internal DNS server and an outgoing SMTP relay, all located in the "Internal internet services" network. E-mails coming from the Internet are received at the external SMTP relay, sitting on the "Screened Subnet" network and then passed onto the internal mail server located at the "Internal Servers" network, where they are accessed by the users via POP3S and IMAPS. Outgoing mail messages are sent from the internal mail servers and routed through the outgoing SMTP relay ("int. SMTP relay") on the Internal Internet Services network.

Employees connected to the branch offices' internal networks have the same needs and these are served in the very same way: their traffic is routed through a VPN to headquarters where it is routed and filtered in the same way as if coming from the "HQ Internal" network. This configuration was chosen because despite the extra cost in network bandwidth and routing resources, it allowed for a centrally managed and better controlled environment.

Currently, there is no need for direct communication between users in different internal networks (HQ and branch offices) and therefore any traffic between offices other than the afore mentioned is denied. If some specific communication need is recognized in the future, the ACLs of the filtering elements would have to be revised to allow that specific traffic.

Source	Destination	Port(s)/Protocol	Description
Employees on internal networks	Mail Server	995/TCP (POP3S) 993/TCP (IMAPS) 25/TCP (SMTP)	Employees access to Mail Server, both to send and receive mail.
Employees on internal networks	File Server	445/TCP (CIFS)	Employees access to File Server.
Employees on internal networks	Application Server	40840/TCP (proprietary)	Employees access to Application Server.
Employees on internal networks	Web and FTP Proxy Server	8088/TCP (web and FTP proxy)	Employees access to Web and FTP Proxy Server
Web and FTP Proxy Server	Internet	80/TCP (WWW) 443/TCP (HTTPS) 21/TCP (FTP) 20/TCP (FTP)	Web and FTP access to the Internet from the Web and FTP Proxy Server.
Employees on internal networks	Internal DNS Server	53/UDP (DNS)	Employees access to Internal DNS Server.
All servers except on screened subnet	Internal DNS Server	53/UDP (DNS)	All GIAC servers except those on the screened subnet also need access to the Internal DNS Server. Servers on the screened subnet will use the DNS servers of the ISP.
Internal DNS Server	Internet	53/UDP (DNS)	DNS access to the Internet from the Internal DNS Server
Mail Server	Outgoing Mail Relay	25/TCP (SMTP)	Mail Server access to Outgoing Mail Relay in order to send mail to the Internet.
Outgoing Mail Relay	Internet	25/TCP (SMTP)	SMTP access to the Internet from Outgoing Mail Relay.

Table 4 Access requirements for employees on the internal networks (including branch offices)

2.2.5 Remote users (sales force)

Salespeople need to log into the network from remote locations and still have the same connectivity as employees on the internal network

This is achieved using VPNs based on IPSec and L2TP.

Thus, the access requirements are the same as those of the employees on the internal network plus the access to the VPN concentrators[17] in order to establish the VPN. When they set up the VPN, they are assigned an IP address in a virtual network with the same access privileges as the internal network.

Source	Destination	Port(s)/Protocol	Description
Salespeople (anywhere on the Internet)	External firewall	500/UDP (IKE) IP protocol 50 (ESP)	Remote users (salespeople) access to establish a VPN with headquarters.
Note: *Additionally, the same access requirements of employees on the internal network also apply here to salespeople.*			

Table 5 Access requirements for remote users (salespeople)

2.2.6 System and network administrators

A special group of employees are the system and network administrators. They need SSH access to all systems and network devices to remotely administer them.

They need this access both locally, when they are sitting in the headquarters office, and remotely, when they are on call.

Additionally, on emergency situations they may need unlimited IP access to everywhere.

In order to distinguish between them and regular users, a specific network (named "Support") in the headquarters office is dedicated to them. Physical access to this network is restricted.

Also, when they connect remotely to the network via VPN, if they authenticate themselves successfully as administrators they are assigned an IP address from a specific virtual network (named "Remote Users – Support").

17 In this case the external firewall assumes the function of VPN concentrator as well.

Source	Destination	Port(s)/ Protocol	Description
Administrators (locally and remotely)	Any	22/TCP (SSH)	Administrators access to all systems and network devices.
Administrators (locally and remotely)	Any	Any	Administrators unlimited access to all systems and network devices. **THIS RULE SHOULD BE ACTIVATED IN CASE OF EMERGENCY ONLY AND DEACTIVATED AGAIN AS SOON AS THE EMERGENCY IS OVER**
Note: *Additionally, the same access requirements of* ***employees on the internal network*** *and* ***remote users*** *also apply here to administrators.*			

Table 6 Access requirements for system and network administrators.

2.2.7 General Public

The general public needs access to GIAC's public web server, named "Web Server for General Public" or "WPub" for short), using protocols HTTP and HTTPS, and to GIAC's external mail relay (SMTP) in order to send mail to GIAC users.

They also need access to an external DNS server that translates names into IP addresses for the "giac.com" and "giacentreprises.com" domains, but this function is delegated to the DNS server of the ISP and therefore this traffic does not need to be regulated by GIAC.

Source	Destination	Port(s)/Protocol	Description
Any	WPub	80/TCP (HTTP) 443/TCP (HTTPS)	General public access to the public web server.
Any	External SMTP Relay	25/TCP (SMTP)	General public access to the external mail relay.

Table 7 Access requirements for system and network administrators.

2.3 Network Diagram and IP Addressing Scheme

2.3.1 Network Diagram

Network Diagram

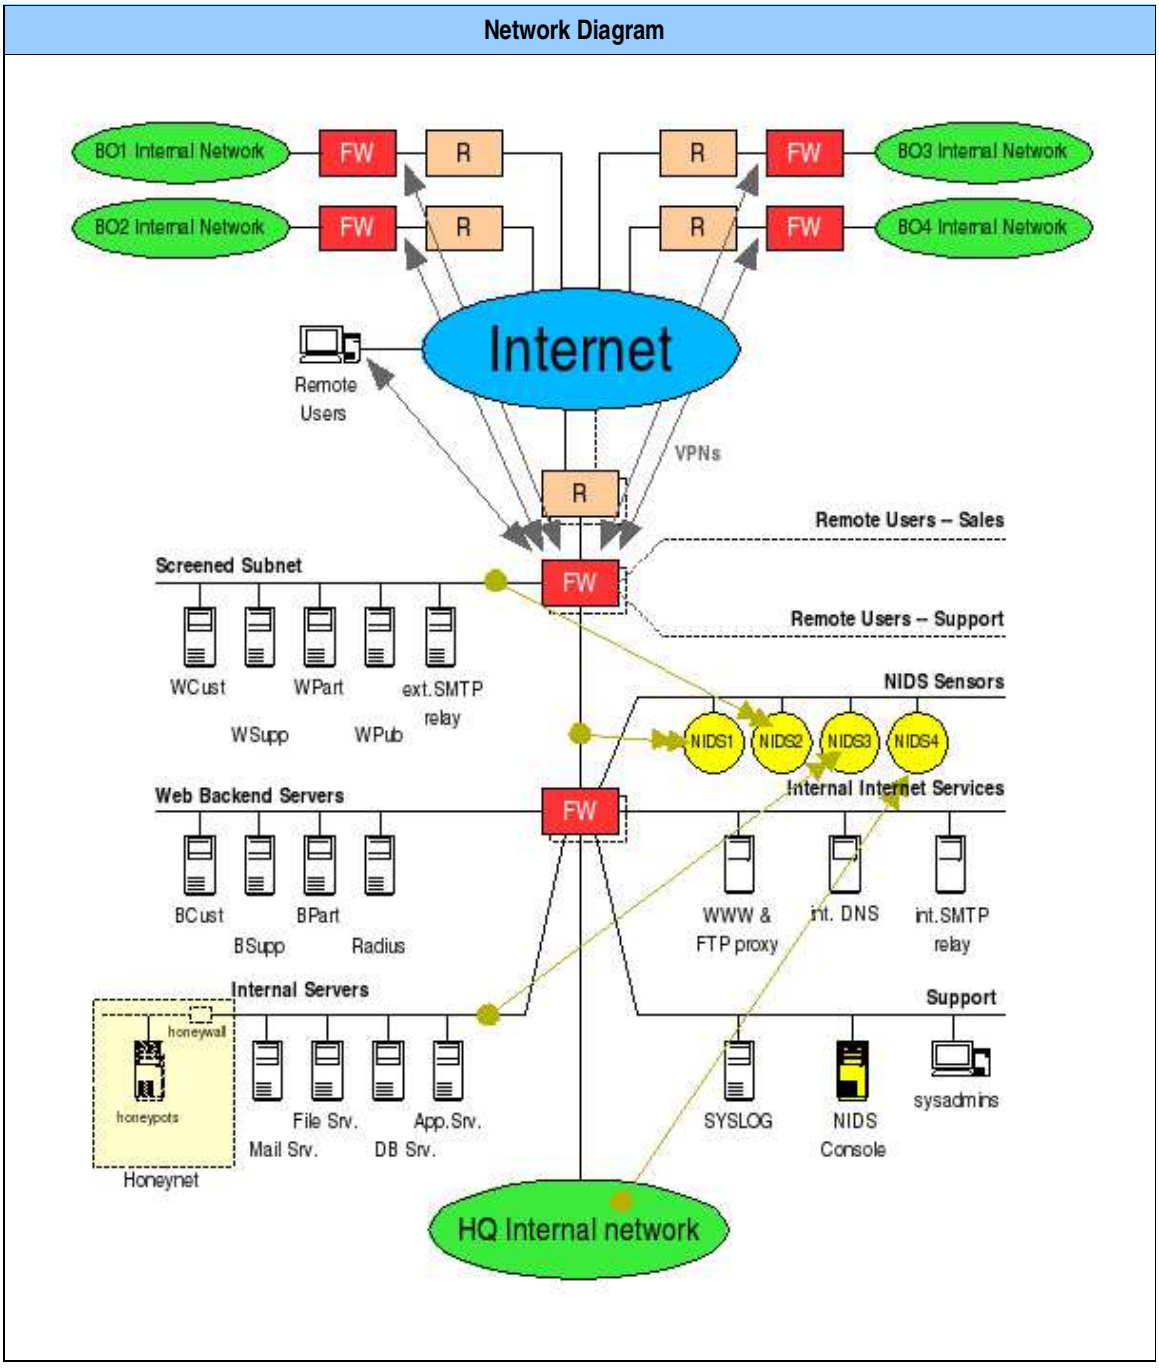

Table 8 Network diagram.

2.3.2 IP Addressing Scheme

GIAC owns five distinct public IP address blocks:

- H1.H2.H3.0/24
- A1.A2.A3.A4/30
- B1.B2.B3.B4/30
- C1.C2.C3.C4/30

The first block is a whole class C network (255 IPs) and is assigned to their headquarters office. GIAC currently does not need that many addresses it has acquired them as a precaution for future growth.

The other four blocks are only 4 IPs big each and correspond to the 4 remote offices GIAC has around the world. The remote offices do not need more IPs because they don't offer any service to the Internet. All of their communication with the external world goes through VPN tunnels to headquarters and from there it is forwarded appropriately. They only use 2 IP public IP addresses at each branch office: one for the LAN interface of the border router and the other for the external interface of the firewall.

The class C network has been subneted to accommodate two subnets: a very small subnet linking the border routers and the external firewalls (H1.H2.H3.0/29) and a little bigger subnet holding the externally facing servers (H1.H2.H3.32/28). This scheme allows for future growth of those subnets by simply changing the network mask and also allows for many other subnets to be created in the future as needed.

Only the above networks, destined to be directly accessed from the Internet, use public IP addresses. All the other subnets in GIAC's network make use of the private address space defined in RFC1918. Namely, all subnets are C class networks in the 192.168/16 prefix.

For all communications going from internal systems (with private addresses) to the Internet, NAT (network address translation) is performed at the external firewall of headquarters.

Table 9 shows the IP addressing scheme in full detail.

Network name	Network address	Description
HEADQUARTERS		
Link between border routers and ext. firewalls	*H1.H2.H3.0 / 29*	Small network linking the border routers with the external firewalls.
External Routers (floating IP)	*H1.H2.H3.1*	
External Firewalls (floating IP)	*H1.H2.H3.2*	
Link between ext. firewalls and int. firewalls	*H1.H2.H3.16 / 29*	Small network linking the external firewalls with the internal firewalls.
External Firewalls (floating IP)	*H1.H2.H3.17*	
Internal Firewalls (floating IP)	*H1.H2.H3.18*	
Screened Subnet	*H1.H2.H3.32 / 28*	Network holding externally facing servers: external SMTP relay, WPub, WCust, WSupp, WPart.
External Firewalls (floating IP)	*H1.H2.H3.33*	
WCust	*H1.H2.H3.36*	
WSupp	*H1.H2.H3.37*	
WPart	*H1.H2.H3.38*	
WPub	*H1.H2.H3.39*	
Ext. STMP Relay	*H1.H2.H3.40*	
Remote Users - Sales	192.168.10.0 / 24	Virtual network for remote salespeople connecting via VPN.
External Firewalls (floating IP)	*192.168.10.1*	
Remote Users - Support	192.168.11.0 / 24	Virtual network for remote support personnel (system and network admins) connecting via VPN.
External Firewalls (floating IP)	*192.168.11.1*	
NIDS Sensors	192.168.12.0 / 24	Network holding the IP enabled interface of NIDS sensors. Their other interface is in listen-only mode, without IP address, connected to switch ports in monitor mode.
Internal Firewalls (floating IP)	192.168.12.1	
Internal Internet Services	192.168.13.0 / 24	Network holding servers for common Internet services for the internal network.
Internal Firewalls (floating IP)	192.168.13.1	
WWW and FTP Proxy	192.168.13.10	
Int. DNS	192.168.13.11	
Int. SMTP Relay	192.168.13.12	
Web Backend Servers	192.168.14.0 / 24	Network holding the backend application servers for serving customers, suppliers and partners.
Internal Firewalls (floating IP)	192.168.14.1	

Network name	Network address	Description
BCust	192.168.14.10	
BSupp	192.168.14.11	
BPart	192.168.14.12	
Radius	192.168.14.13	
Internal Servers	192.168.15.0 / 24	Network holding internal servers (e.g.: mail, file, database and application servers).
Internal Firewalls (floating IP)	192.168.15.1	
Mail Server	192.168.15.10	
File Server	192.168.15.11	
DB Server	192.168.15.12	
App Server	192.168.15.13	
Support	192.168.16.0 / 24	Network dedicated to support personnel (sysadmins, netadmins, operators).
Internal Firewalls (floating IP)	192.168.16.1	
Syslog	192.168.16.10	
NIDS Console	192.168.16.11	
Workstations of System Administrators	192.168.16.12	
HQ Internal Network	192.168.64.0 / 24	Internal network containing all regular users.
Internal Firewalls (floating IP)	192.168.64.1	
BRANCH OFFICE 1 (BO1)		
BO1 External	*A1.A2.A3.A4 / 30*	Link between border router and firewall.
Border Router	*A1.A2.A3.A4+1*	
Firewall	*A1.A2.A3.A4+2*	
BO1 Internal	192.168.1.0 / 24	BO1 internal network.
Firewall	192.168.1.1	
BRANCH OFFICE 2 (BO2)		
BO1 External	*B1.B2.B3.B4 / 30*	Link between border router and firewall.
Border Router	*B1.B2.B3.B4+1*	
Firewall	*B1.B2.B3.B4+2*	
BO1 Internal	192.168.2.0 / 24	BO2 internal network.
Firewall	192.168.2.1	
BRANCH OFFICE 3 (BO3)		
BO1 External	*C1.C2.C3.C4 / 30*	Link between border router and firewall.
Border Router	*C1.C2.C3.C4+1*	

Network name	Network address	Description
Firewall	*C1.C2.C3.C4+2*	
BO1 Internal	192.168.3.0 / 24	BO3 internal network.
Firewall	192.168.3.1	
BRANCH OFFICE 4 (BO4)		
BO1 External	*D1.D2.D3.D4 / 30*	Link between border router and firewall.
Border Router	*D1.D2.D3.D4+1*	
Firewall	*D1.D2.D3.D4+2*	
BO1 Internal	192.168.4.0 / 24	BO4 internal network.
Firewall	192.168.4.1	

Table 9 IP Addressing scheme

2.4 Architecture Components

2.4.1 Filtering routers

The filtering routers are the border routers connecting the headquarters and the branch offices networks to the Internet.

The border router in all four branch offices is a Cisco 1760 Modular Access Router. At headquarters there are two Cisco 2651XM Multiservice routers, each of them connected to a different ISP, and configured in high availability using HSRP. They all run version 12.3(10) of Cisco's operating system IOS.

Since all GIAC business is conducted online, loosing Internet connectivity is a risk that must be mitigated at (almost) all costs. That is the reason why the border routers are duplicated and connected to different ISPs.

They constitute the first barrier against unwanted traffic. Consequently, ACLs both ingress and egress have been configured on them. They block the following traffic (in both directions except where noted otherwise):

- Packets claiming to come from or be destined to the private non-routable address ranges defined in RFC1918[18].
- Packets claiming to come from or be destined to the loopback network (127.0.0.0/8) defined

18 NAT and VPNs are performed at the firewalls. Therefore, border routers should never see a packet with a private address although GIAC uses private addressing internally.

in RFC1700.

- Packets going into one interface claiming to come from a network accessible through a different interface according to the routing table (antispoofing).
- Any traffic directed to the router itself with the only exception of NTP from specific time servers (to keep clock synchronization) and SSHv2 from specific IP addresses (to allow remote administration).
- Source routed packets.
- ICMP redirect packets.

Additionally, all routers have been hardened so that they do not run any unnecessary network services.

2.4.2 Firewalls

Each branch office has a Nokia IP260 firewall appliance [NOK01] running Checkpoint VPN-1 over IPSO 3.81 [NOK02] behind the border router. At headquarters, behind the border router there is a cluster of two Nokia IP530 firewall appliances, configured in high availability, running Checkpoint VPN-1 over IPSO 3.81. These will be called the external firewalls or simply the external firewall. Behind them, there is another cluster of firewalls, this time Cisco PIX 535 systems running Cisco PIX Firewall Software Version 6.3 [CIS01]. These will be called the internal firewalls or simply the internal firewall.

The internal firewalls were selected to be from a different manufacturer in all respects (hardware, operating system and firewall software) than the external firewalls to mitigate the risk of an intrusion caused by a vulnerability found on the firewalls. The rationale is that the probability of a vulnerability being discovered in both vendors' products at the same time is close to zero. If a vulnerability were discovered in the Checkpoint firewall software that allowed attackers to get full control over it, an attacker trying to get into the internal networks of GIAC would still need to overcome the limitations imposed by the internal Cisco PIX.

Both the external and internal firewalls at headquarters are duplicated for redundancy (high availability). If one of them fails, its counterpart takes over and connectivity is maintained. Again, since all GIAC business is conducted online, loosing Internet connectivity is a risk that must be mitigated at (almost) all costs. Replicating the firewalls and having border connected to 2 different ISPs ensures that a single failure will not kill GIAC's Internet connectivity. Some other elements which are not replicated may fail, like a web server, but this would cause only a partial damage as opposed to total isolation.

The firewalls are key components in enforcing the security policy. Thus, a tight configuration and proper maintenance is critical. They allow through them only the traffic explicitly permitted in the policy, disallowing anything else. This configuration ("deny by default") is far more secure than the opposite ("allow by default") because it doesn't leave any space to an attacker to exploit vulnerabilities in non-essential services.

2.4.3 VPNs

There are two different types of VPNs in GIAC Enterprises network. The first type is the router-to-router permanent (actually "firewall-to-firewall" in this case) VPN established between each branch office and headquarters. The second is the host-to-router (actually "host-to-firewall" in this case) VPN established between remote users' laptops and headquarters.

All traffic leaving the branch offices, even traffic intented to the Internet, like web browsing is first routed to headquarters via VPN and then it is filtered as if coming from the headquarters internal network.

The VPNs between branch offices and headquarters are IPSec ESP tunnels that encapsulate and encrypt all traffic going from the branch office to headquarters and vice versa, without modifying the private addressing of the packets. The endpoints of the tunnel for each of the VPNs are the firewall of the branch office on one side and the external firewall of headquarters on the other side.

The endpoints could have been chosen to be different. For example, the VPN could have been configured between the border routers or between a border router and a firewall. Or dedicated VPN gateways could have been used instead. The difference between the different options in terms of security is not very high, but the current configuration was chosen because it offered some small advantages. The use of dedicated VPN gateways was discarded because given the not too high network load that could easily be managed by the firewalls or routers the cost of buying extra equipment would not be justified. The firewalls were selected as the endpoints over the routers or mixed options because it allowed GIAC to double check the encryption was taking place properly by simply sniffing the ethernet link between the firewalls and the border routers. Should the encryption take place at the routers, the only way to verify the encryption would be to relay on the router log. On top of that, because NAT was also chosen to be performed at the external firewalls[19], tunneling the VPN traffic right from them meant that the firewalls would be the frontier for privately addressed traffic.

Authentication of both ends is performed using digital certificates generated by GIAC's own certification authority (CA) which is nothing more than a Windows 2003 Server Enterprise Edition.

19 NAT is performed at headquarters' external firewalls. At branch offices, NAT is not required since all traffic is sent to headquarters via VPN tunnels retaining the private addressing.

See section “Additional Components” later for more information on the CA.

The second type of VPN (host-to-router) is used by remote users to log into the headquarters network. Remote users use the client VPN application “Checkpoint SecureClient VPN”, by Checkpoint, in their laptops to establish a VPN with headquarters external firewalls. Again, both ends (computers) are authenticated using digital certificates and the user authenticates him or herself using username and password.

Once authenticated, users receive a dynamic IP address belonging to a virtual network defined in the firewall configuration, and all traffic from then on is routed through the VPN to headquarters: no split tunneling allowed. Split tunneling happens when some traffic is routed through the VPN while other traffic is sent via some other route (e.g. directly to the Internet). Not allowing split tunneling ensures that all traffic leaving the box will be encrypted and pass all filters set up at the other end of the tunnel, in this case the headquarters filtering rules.

There are two types of remote users, each with different access requirements, namely salespeople and system (or network) administrators. A different virtual network was defined for each group so that when users dial in they get an IP address belonging to their appropriate network and so their access requirements can be fulfilled and enforced according to their IP network address.

2.4.4 Network based IDS sensors

Prevention is important, but early detection of attacks is just as important. In order to detect network based attacks and act on them as soon as possible, GIAC has deployed a number of network based intrusion detection systems (NIDS) across the network.

The NIDS solution selected is composed of several Sourcefire Intrusion Sensors and the central console called Sourcefire Defense Center. The intrusion sensors monitor the activity of the network they are attached to looking for a specific set of malicious traffic signatures and when they detect some suspicious traffic they send this information to the central console where an alert is generated.

Initially, to keep budget under control only four sensors were deployed to the four most critical spots of the network.

One of them monitors the traffic of the screened subnet where the externally facing servers are located. Since this is a network exposed to Internet a high number of unsuccessful attacks is expected. In order to keep a low level of noise this sensor is configured to log all attack attempts but only alert on events that clearly denote a successful attack, like a known successful attack response or anomalous traffic generated by one of the servers.

The second sensor monitors the traffic between the external and the internal firewalls. This

sensor would catch incoming or outgoing attacks not stopped by the firewalls. The expected number of attacks crossing this link, both outwards and inwards, is close to zero. Therefore, the sensor is configured to alert on any attack it sees hoping the number of false positives will be very small.

A third sensor monitors the internal servers network and is set to alert on any attack attempt against the servers or any suspicious response. Attacks coming from the Internet are very unlikely here since they would have to go past two firewall layers and internal attacks would still have to cross the internal firewall, which makes it also difficult to an attacker to reach the servers, but the possibility exists and this is were the jewels crown is held, so better err on the safe side.

Finally, a fourth sensor monitors the internal network. This sensor is expected to alert on attacks coming from and/or going to internal users. The criticality of these attacks is probably lower than the other attacks, but having an early warning if an internal system gets infected by some virus or worm and starts scanning other systems may save a lot of time an money in recovering from such outbreak.

All sensors are located in a dedicated network ("NIDS Sensors") from which they feed their data to the NIDS console in the Support network, through the internal firewall which allows only the exact ports required for this communication. Each sensor has a second ethernet adapter in listen-only mode connected to the appropriate switch port in mirror mode, thus monitoring the appropriate network segment. This isolation of the NIDS sensors is necessary in case an attacker managed to executed arbitrary code on them. Although they are not directly accessible via IP on the monitoring interface, it has already been the case that a particularly crafted packet traveling the network and being passively sniffed by a NIDS sensor could cause the execution of arbitrary code in the sensor because of some vulnerability in the packet processing software [ISS01]. Isolating them in a tightly controlled network drastically reduces the chances of an attacker being able to take advantage of such a flaw.

The central console is located in the Support network where is to be accessed by authorized administrators only.

2.4.5 Additional components

A few additional components are included in GIACs security architecture. These are commented below.

Syslog Server

Good logs are an indispensable tool for incident handlers. GIAC has centralized all logs from critical systems and network devices in a single syslog server located in the Support network. A simple but powerful PC running GNU/Linux Fedora Core 3 does the job.

This, together with a good time synchronization using NTP provides support personnel with a wealth of valuable information to detect and investigate security and non-security incidents.

Certification Authority (CA)

GIAC uses digital certificates for various purposes. Some of those certificates were issued by Verisign while others were generated using a root CA set up internally by GIAC.

Certificates for the externally facing web servers were bought to Verisign Inc. so that they could be recognized and verified instantly by everybody on the Internet.

However, it was decided that digital certificates to be used internally would be generated using the CA capabilities included in Microsoft Windows Server 2003. The decision was mainly cost driven: as the number of remote employees and remote offices grow the number of required certificates will also grow; besides, the Windows Server was already in-house for other applications and thus its CA functionality could be used at no extra cost.

It is recommended to keep any root CA offline for security reasons, but GIAC decided to use the same server that was online, on the Internal Servers network, offering some other services. This decision was taken after weighting the risks against the benefits. The benefits of having it online were clear: a spare server (hardware) and an extra OS license would be needed to have it offline and that represented a non-negligible cost. On the other hand, the risk being mitigated was that if the CA was compromised then all certificates generated by it or any subordinated CA (none in this case) would have to be replaced. This risk seemed acceptable given the small number of certificates that would be needed initially. Should the company grow much bigger and therefore need many more certificates, this risk should be re-evaluated and probably an extra offline CA would then be found cost efective.

Antivirus and Personal Firewall Software

Antivirus software is used in the incoming mail server and the web and ftp proxy servers to reduce the probability of malware getting into the network. Additionally, antivirus and personal firewall software is used on each user's workstation to avoid infection of these systems, which could cause a virus or worm outbreak in the internal network.

It may be argued if these are “perimeter” defenses, but what cannot be questioned is that these are real and vital defenses in any network nowadays.

Honeynet

Last, but not least, a honeynet is included in GIAC's security architecture.

The honeynet is configured as part of the Internal Servers network. Its objective is to catch and

report any connection attempt to the unused IP address space of that network. Such connection attempts would reveal a certainly unwanted and potentially malicious scanning activity in a sensitive zone as the Internal Servers network is.

The internal firewall is configured to block all traffic directed to the real servers except on those specific ports where approved services are running, and to allow any traffic directed to any other IP addresses on that network.

The honeynet is built using a simple PC with a single network interface, VMWare workstation 4.1 and two virtual machines configured: a Honeywall, booting off the Honeywall CD image from The Honeynet Project [HON03], and a GNU/Linux Fedora Core 3 system running "honeyd". The Honeywall is configured to only allow incoming traffic and responses, never connections initiating on the honeynet, to prevent the possibility of an attacker using the honeypots to launch new attacks. Honeyd is configured to emulate systems with the same personality and services as the real servers.

Actually, any scanning activity against this network segment could be detected simply by looking at the log of the internal firewalls. However, having the honeynet in place allows the security team to gather extra information from those scans or attacks. In the firewall log only the source and destination IPs and ports would be noted. The honeynet, allowing connections to get established with the emulated services can capture samples of malicious payload should the attacker care to send it against the honeypots.

2.5 Implementing Defense in Depth

The concept of Defense in Depth refers to establishing defense controls or countermeasures in layers so that if one control is defeated by attackers there will still be other controls protecting the valuable information they are after. Defense in Depth is the opposite to an all-or-nothing approach where the valuables are protected by one or many countermeasures in such a way that as soon as the attackers overcome one of such defense controls they gain access to all the valuables. In general, the more the layers, the more complicated gets for attacker to achieve their objectives and therefore the better.

GIAC's security architecture is built around this concept. To begin with, the internal servers, which, contain the crown jewels (the data stored in them), are located in a very protected network ("Internal Servers"). Attackers from the Internet will have to go over the filtering imposed by the border router, the filtering imposed by the external firewalls and the filtering imposed by the internal firewalls.

Being the firewalls from different vendors and having many redundant filtering rules in all filtering devices, including the border routers, makes it almost impossible that a single vulnerability that could be found in any of these filtering products would open all the way at once for attackers to get

to the end. The probability of the same vulnerability affecting all of them is very close to zero and the probability of different vulnerabilities of the different products being found at the same time is again very low, although you can never say never.

With the filtering devices in place, direct connection from the Internet is allowed only to specific services on specific servers on the Screened Subnet. Getting ahold of one of those servers would probably be the first stage in an external attack. But these servers have been hardened, so that their configuration is as restrictive as possible, and a patching policy is in place to keep them (specially the externally accessible services) always up to date, that is, without known vulnerabilities. Yet new vulnerabilities can always be discovered by attackers before patches are available so the compromise of these systems is always a possibility. The network intrusion sensor installed in this network helps mitigate this risk: the attack would not be stopped short by the sensor but at least it will probably alert on it allowing humans to take action and kick the attackers off.

Even if the compromise of such system went undetected, attackers would still have to find another vulnerability in the few systems accessible from it, namely the backend servers or the internal DNS server, and they would have to find such vulnerability in the few services that are allowed through the external and the internal firewalls.

And so the story goes on.

Hardened systems, locked down services, filtering rules, intrusion detection sensors, encryption (SSL and VPNs), the honeynet, extense logging capabilities and proper operation and response to incidents all add together to make a successful attack a task as difficult as possible.

3 Router and Firewall Policies

3.1 General Security Stance

The mission of the border router and the external firewall at GIAC headquarters is to move allowed traffic to and from GIAC's network as fast as possible while preventing any unwanted traffic to enter or leave GIAC's network.

Both elements apply filtering rules to discard unwanted traffic. Some rules are complementary between the two devices but most of them are redundant thus providing defense in depth.

Filtering is applied both to inbound and outbound traffic. The reason to filter inbound traffic is obvious: it is important to prevent potentially malicious traffic from entering the network in order to reduce the chances of a system being compromised. The reason to filter outbound traffic is actually twofold: being a good Internet neighbor by preventing potentially malicious traffic from leaving GIAC's network is one side of it, the other side is making it harder for a backdoor, that an attacker somehow manages to install inside the network, to “call home”.

Both filtering rulesets are built with a “deny by default” approach: everything not specifically allowed is denied. This makes for a much more restrictive posture than the opposite, “allow by default”, where anything not specifically denied would be allowed.

Although there are two border routers and two external firewalls there is only one filtering policy for each type of device. That's because being configured as high availability clusters both members of the cluster share and enforce the same policy.

3.2 HQ Border Router Policy

Filtering rules in Cisco routers are called access control lists (ACL). Exactly one ACL can be applied to each interface for inbound traffic (traffic entering the router through that interface) and another for outbound traffic (traffic leaving the router through that interface). If no ACL is applied to a specific direction on a specific interface, a default “allow all” ACL is assumed.

GIAC has decided to apply strict ACLs to all traffic as soon as it enters the router, both from the Internet through the Serial0 (WAN) interface (ingress filtering) and from GIAC's network through the Ethernet0 (LAN) interface (egress filtering). This follows the generally accepted advice to filter traffic as soon as it enters the device instead of just before leaving the device in order to avoid unnecessary CPU cycles routing packets that will later be discarded.

Rule ordering is important. A packet crossing an interface with an ACL applied to it will be checked against each rule in the ACL in order and as soon as it matches the selection criteria of one of them the corresponding action will be performed (allow, deny) and the packet will not be

checked against the remaining rules. This is sometimes referred to as the "first match out" algorithm. Thus, if a packet would match the selection criteria of two rules, one indicating action "allow" and the other indicating action "deny", the packet will be allowed or denied depending on which of the two rules comes first in the rule set.

This algorithm implies that exceptions to specific allow rules must be included in the rule set before those allow rules and that very often matched rules should be set as early as possible in the rule set to improve performance: as soon as a packet is matched against a rule it the matching stops for that packet, thus saving CPU cycles.

Apart from the rules listed below, the command "no ip source-route" is used to deny any source routed packets.

3.2.1 Ingress Filtering

Table 10 lists the ingress filtering rules applied to the "in" (inbound) direction of the Serial0 interface. A detailed explanation of each rule follows.

#	Source	Destination	Port(s)/Protocol	Action
1	10.0.0.0/8 172.16.0.0/12 192.168.0.0/16 127.0.0.0/8	Any	Any	Deny
2	Any	10.0.0.0/8 172.16.0.0/12 192.168.0.0/16 127.0.0.0/8	Any	Deny
3	H1.H2.H3.0/24	Any	Any	Deny
4	Any	H1.H2.H3.0/24	TCP Established	Allow
5	Any	H1.H2.H3.2	500/UDP IP Protocol 50 (ESP)	Allow
6	Any	H1.H2.H3.39	80/TCP (HTTP) 443/TCP (HTTPS)	Allow
7	Any	H1.H2.H3.40	25/TCP (SMTP)	Allow
8	Any	H1.H2.H3.36 H1.H2.H3.37 H1.H2.H3.38	443/TCP (HTTPS)	Allow
9	Any	H1.H2.H3.0/24	UDP	Allow
10	Any	Any	Any	Deny

Table 10 HQ Border Router ingress policy

Rule #1 drops any traffic claiming to come from networks in the private address ranges defined

by RFC1918 or from the loopback network (127.0.0.0/8) defined in RFC1700. Traffic from these networks should never be seen on the Internet. It should not be routed by any Internet router, but some ISPs are more strict than others so it is better not to assume that these packets will never arrive.

Rule #2 does the same for traffic claiming to be destined to those networks.

Rule #3 drops any traffic claiming to come from GIAC's public class C network. Traffic arriving here should always be destined to that network, but never have a source IP address in that range.

Rule #4 accepts any packet belonging to a previously established TCP session and addressed to GIAC's address space. This rule is the first allow rule to improve performance because most packets will match it.

Rule #5 accepts VPN packets destined to the VPN tunnel end: the external firewall. The source address of those packets is “any” because it includes the VPNs from the branch offices, which are fixed IPs, but also the VPNs from remote users around the Internet.

Rules #6 to #8 allow incoming TCP connections to the services offered by the servers on the Screened Subnet.

Rule #9 allows UDP packets destined to GIAC. GIAC does not offer any service via UDP, but needs to receive UDP replies. The firewall will make sure only real replies are accepted in.

Rule #10 drops everything else. It would not be necessary since the default rule when there is an ACL applied is drop everything else, but it is good to include it for clarity.

3.2.2 Egress Filtering

Table 11 lists the egress filtering rules applied to the “in” (inbound) direction of the Ethernet0 interface. A detailed explanation of each rule follows.

#	Source	Destination	Port(s)/Protocol	Action
1	10.0.0.0/8 172.16.0.0/12 192.168.0.0/16 127.0.0.0/8	Any	Any	Deny
2	Any	10.0.0.0/8 172.16.0.0/12 192.168.0.0/16 127.0.0.0/8	Any	Deny
3	H1.H2.H3.0/24	Any	TCP Established	Allow
4	H1.H2.H3.2	H1.H2.H3.0/24	80/TCP (HTTP) 443/TCP (HTTPS) 20/TCP (FTP) 21/TCP (FTP) 25/TCP (SMTP) 123/TCP (NTP) 53/UDP (DNS) IP Protocol 50 (ESP)	Deny
5	H1.H2.H3.2	Any	80/TCP (HTTP) 443/TCP (HTTPS) 20/TCP (FTP) 21/TCP (FTP) 25/TCP (SMTP) 123/TCP (NTP) 53/UDP (DNS) IP Protocol 50 (ESP)	Allow
6	H1.H2.H3.32/28	ISP DNS Servers	53/UDP (DNS)	Allow
7	Partner router fixed IP	224.0.0.2	1985/UDP	Allow
8	H1.H2.H3.2	H1.H2.H3.1	22/TCP (SSH)	Allow
9	Any	Any	Any	Deny

Table 11 HQ Border Router egress policy

Rules #1 and #2 are exactly the same as before: deny packets with private addresses.

Rule #3 allows packets from established TCP connections. It is the same as before but GIAC's address space is now the source instead of the destination.

Rules #4 and #5 allow traffic from GIAC to the approved list of services to anywhere in the Internet except GIAC's own address range, which includes the router itself.

Rule #6 allows DNS queries from the servers on the Screened Subnet to the DNS servers of the ISP.

Rule #7 allows HSRP (Hot Standby Router Protocol) traffic from the partner router. This is needed to have high availability.

Rule #8 allows SSH traffic to the router itself from GIAC (the true source IP will be NATed) to allow for remote administration.

Finally, rule #9 drops everything else mirroring the default rule when an ACL is applied. Again, it is included for clarity.

3.3 HQ External Firewalls Policy

Rule ordering is also important in Checkpoint firewalls as it is in Cisco devices. The same "first match out" algorithm is applied in checking packets against the rule set. See the previous section for more information on the implications of this algorithm.

3.3.1 Filtering Policy

Table 12 lists the filtering rules of the external firewall. A detailed explanation of each rule follows.

#	Source	Destination	Port(s)/Protocol	Action
1	Any	H1.H2.H3.2	500/UDP IP Protocol 50 (ESP)	Accept
2	H1.H2.H3.2	A1.A2.A3.A4+2 B1.B2.B3.B4+2 C1.C2.C3.C4+2 D1.D2.D3.D4+2	500/UDP IP Protocol 50 (ESP)	Accept
3	Any	H1.H2.H3.2	Any	Drop
4	Any	H1.H2.H3.39	80/TCP (HTTP) 443/TCP (HTTPS)	Allow
5	Any	H1.H2.H3.40	25/TCP (SMTP)	Allow
6	Any	H1.H2.H3.36 H1.H2.H3.37 H1.H2.H3.38	443/TCP (HTTPS)	Allow
7	Any	H1.H2.H3.0/24	Any	Drop
8	H1.H2.H3.32/28	ISP DNS Servers	53/UDP (DNS)	Allow
9	192.168.13.10	Any	80/TCP (HTTP) 443/TCP (HTTPS) 20/TCP (FTP) 21/TCP (FTP)	Allow
10	192.168.13.11	Any	53/UDP (DNS)	Allow
11	192.168.13.12	Any	25/TCP (SMTP)	Allow

#	Source	Destination	Port(s)/Protocol	Action
12	192.168.11.0/24 192.168.16.0/24	H1.H2.H3.0/24 A1.A2.A3.A4/30 B1.B2.B3.B4/30 C1.C2.C3.C4/30 D1.D2.D3.D4/30 192.168.0.0/16	22/TCP (SSH)	Allow
13	192.168.1.0/24 192.168.2.0/24 192.168.3.0/24 192.168.4.0/24 192.168.10.0/24 192.168.11.0/24	192.168.15.10	995/TCP (POP3S) 993/TCP (IMAPS) 25/TCP (SMTP)	Allow
14	192.168.1.0/24 192.168.2.0/24 192.168.3.0/24 192.168.4.0/24 192.168.10.0/24 192.168.11.0/24	192.168.15.11	445/TCP (CIFS)	Allow
15	192.168.1.0/24 192.168.2.0/24 192.168.3.0/24 192.168.4.0/24 192.168.10.0/24 192.168.11.0/24	192.168.15.13	40840/TCP (proprietary)	Allow
16	192.168.1.0/24 192.168.2.0/24 192.168.3.0/24 192.168.4.0/24 192.168.10.0/24 192.168.11.0/24	192.168.13.10	8088/TCP (web & FTP proxy)	Allow
17	192.168.1.0/24 192.168.2.0/24 192.168.3.0/24 192.168.4.0/24 192.168.10.0/24 192.168.11.0/24	192.168.13.11	53/UDP (DNS)	Allow
18	H1.H2.H3.36	192.168.14.10	443/TCP (HTTPS)	
19	H1.H2.H3.37	192.168.14.11	443/TCP (HTTPS)	
20	H1.H2.H3.38	192.168.14.12	443/TCP (HTTPS)	
21	H1.H2.H3.36 H1.H2.H3.37 H1.H2.H3.38	192.168.14.13	1812/TCP (RADIUS) 1812/UDP (RADIUS) 1813/TCP (RADIUS ACCT) 1813/UDP (RADIUS ACCT)	
22	Any	Any	Any	Deny

Table 12 HQ External Firewalls filter policy

Rule #1 allows VPN traffic directed to the firewall itself, which is the tunnel terminator.

Rule #2 allows VPN traffic sent from the firewall to the branch offices.

The second rule is needed because VPN traffic can be originated in either side, BO or HQ. In the first rule the source IP is any because this rule also includes the VPN traffic from remote (moving) users.

Rule #3 drops any other connection attempts to the firewall itself. Replies to already established connections are allowed implicitly.

Rules #4 to #6 allow traffic to the specific services offered by GIAC servers on the screened subnet.

Rule #7 drops any other traffic directed to GIAC. Again, replies to outgoing connections are implicitly allowed.

Rule #8 allows DNS traffic from the servers on the screened subnet to the DNS servers of the ISP. Once again, replies are implicitly allowed.

Rule #9 allows traffic from the Web and FTP proxy to the Internet. Its source IP will be masquerade by the NAT rules explained later.

Rule #10 allows DNS queries from the internal DNS server to anywhere in the Internet. Replies are implicitly allowed.

Rule #11 allows SMTP traffic from the internal outgoing mail relay to connect to SMTP servers anywhere on the Internet.

Rule #12 allows SSH connections from administrators to all of GIAC addresses.

Rules #13 to #17 allow connections from remote internal networks and from the two virtual networks of remote users (salespeople and administrators) to the appropriate internal servers.

Rules #18 to #21 allow connections from the external servers to their respective backend servers and to the RADIUS server.

Finally, rule #22 drops any other traffic.

Additionally, the “antispoofing” capability is set in the firewall so that packets claiming to come from a network through an interface that is not the intended route to that network are dropped.

3.3.2 NAT Policy

Table 13 lists the NAT rules of the external firewall. A detailed explanation of each rule follows.

#	Source	Destination	Port(s)/Protocol	Action
1	192.168.13.10	Any	80/TCP (HTTP) 443/TCP (HTTPS) 20/TCP (FTP) 21/TCP (FTP)	Hide – external (H1.H2.H3.2)
2	192.168.13.11	Any	53/UDP (DNS)	Hide – external (H1.H2.H3.2)
3	192.168.13.12	Any	25/TCP (SMTP)	Hide – external (H1.H2.H3.2)
4	192.168.16.0/24 192.168.11.0/24	H1.H2.H3.32/28	22/TCP (SSH)	Hide – screened (H1.H2.H3.33)
5	192.168.16.0/24 192.168.11.0/24	H1.H2.H3.0/29 A1.A2.A3.A4/30 B1.B2.B3.B4/30 C1.C2.C3.C4/30 D1.D2.D3.D4/30	22/TCP (SSH)	Hide – external (H1.H2.H3.2)

Table 13 HQ External Firewalls NAT policy

Rules #1 to #4 masquerade the source internal IP address of the internal internet services servers connecting to the Internet using the IP of the external interface of the firewall.

Rule #4 masquerades the source internal IP address of the administrators' SSH connections to servers on the screened subnet using the IP of the interface of the firewall on the screened subnet.

Rule #5 masquerades the source internal IP address of the administrators' SSH connections to remote public GIAC addresses using the IP of the interface of the external interface of the firewall.

3.3.3 VPN Policy

Table 14 lists the VPN rules of the external firewall. A detailed explanation of each rule follows.

#	Source	Destination	Port(s)/Protocol	Action
1	192.168.10.0/24 192.168.11.0/24 192.168.13.0/24 192.168.15.0/24 192.168.16.0/24 192.168.64.0/24	192.168.1.0/24	Any	Encrypt VPN 1 (A1.A2.A3.A4+2)
2	192.168.10.0/24 192.168.11.0/24 192.168.13.0/24 192.168.15.0/24 192.168.16.0/24 192.168.64.0/24	192.168.2.0/24	Any	Encrypt VPN 2 (B1.B2.B3.B4+2)
3	192.168.10.0/24 192.168.11.0/24 192.168.13.0/24 192.168.15.0/24 192.168.16.0/24 192.168.64.0/24	192.168.3.0/24	Any	Encrypt VPN 3 (C1.C2.C3.C4+2)
4	192.168.10.0/24 192.168.11.0/24 192.168.13.0/24 192.168.15.0/24 192.168.16.0/24 192.168.64.0/24	192.168.4.0/24	Any	Encrypt VPN 4 (D1.D2.D3.D4+2)

Table 14 HQ External Firewalls VPN policy

Rule #1 establishes that traffic from the selected internal HQ networks going to branch office number 1 (BO1) must be encrypted using an IPSec tunnel which other end is the firewall at BO1 (A1.A2.A3.A4+2).

Rules #2 to #4 do the same for the other three branch offices.

4 References

[CIS01] Cisco Systems. *Cisco PIX 500 Series Firewalls.* http://www.cisco.com/warp/public/cc/pd/fw/sqfw500/

[CKP01] Checkpoint. *Checkpoint VPN-1.* http://www.checkpoint.com/products/vpn-1_pro/

[FYO01] Fyodor. *Nmap Security Scanner*. http://www.insecure.org/nmap/index.html

[HON01] The Honeynet Project. *Know Your Enemy: Honeynets.* http://www.honeynet.org/papers/honeynet/index.html

[HON02] The Honeynet Project. *Know Your Enemy, 2nd Edition*. http://www.honeynet.org

[HON03] The Honeynet Project. *Honeywall CDROM*. http://www.honeynet.org/tools/cdrom/

[ISS01] Internet Security Systems Inc. *Snort fragmented RPC preprocessor buffer overflow.* http://xforce.iss.net/xforce/xfdb/10956

[KFS01] KeyFocus Ltd. KFSensor: *Practical Windows Honeypot Technology.* http://www.keyfocus.net/kfsensor/

[LIS01] Liston, Tom. *Labrea – Version 2.3 available.* http://www.dshield.org/pipermail/intrusions/2001-November/002301.php

[NET01] Netsec. *Specter Intrusion Detection System*. http://www.specter.com/default50.htm

[NOK01] NOKIA Corp. *Nokia IP260.* http://www.nokia.com/BaseProject/Sites/NOKIA_MAIN_18022/CDA/Categories/Business/NetworkSecurity/IPSecurityPlatforms/DistributedEnterprises/_Content/_Static_Files/nokiaip260-ip265_datasheet_emea.pdf

[NOK02] Nokia Corp. *IPSO*. http://www.nokia.com/cda1?id=46230

[PER01] Perez, David. *Safe at Home?.* http://www.giac.org/practical/GCFA/David_Perez_GCFA.pdf

[PRO01] Provos, Niels. *Honeyd Frequently Asked Questions*. http://www.honeyd.org/faq.php

[SPZ01] Spitzner, Lance. *SANS Wednesday Webcast: Honeypots*. December 01, 2004. https://www.sans.org/webcasts/show.php?webcastid=90525

[SPZ02] Spitzner, Lance. *Honeypot Definition – Almost There.*

http://www.securityfocus.com/archive/119/322363/2003-05-18/2003-05-24/0

[SPZ03] Spitzner, Lance. *Honeytokens: The Other Honeypot.* July 2003. http://www.securityfocus.com/infocus/1713

[STE01] Stevens, Richard W. *TCP/IP Illustrated, Volume 1*. Addison-Wesley, 1994. ISBN 0201633469.

[SYM01] Symantec Corporation. *Symantec Decoy Server.* http://enterprisesecurity.symantec.com/products/products.cfm?ProductID=157

[XPR01] Yarochkin, Fyodor; Arkin, Ofir. *Xprobe.* http://www.sys-security.com/html/projects/X.html

[ZAL01] Zalewski, Michal. *p0f.* http://lcamtuf.coredump.cx/p0f.shtml