



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Ethical Deception and Preemptive Deterrence in Network Security

Brian McFarland
GCFW Practical
Version 4.1

Date: Apr. 11, 05

Table of Contents

Assignment 1: Future State of Security Technology	1
Abstract	1
Background / Introduction	2
Defining Honeypots	2
Tarpits	4
The Honeynet Project	5
Ethics of Honeypot Use	6
Legality of Honeypots	8
Impact on Perimeter Security	9
Conclusion	11
Assignment 2: Security Architecture	12
Access and Security Requirements	12
Customers	12
Suppliers	12
Partners	13
Employees	14
Sales/Tele-workers	14
General-Public	15
Data Flows	15
General Public	15
Customer	16
Supplier	16
Partner	16
Employees	16
Sales/Teleworkers	17
Architecture Components	18
Filtering Router	18
Firewalls	19
VPN	19
Network based IDS sensor(s)	19
Additional Components	20
Network Diagram and IP Addressing Scheme	20
Implementing Defense in Depth	21
Intrusion Detection	22
Intrusion Prevention	22
Host/Service Hardening	22
Deception and Deterrence	23
Authentication and Authorization	23
Anti-Virus	24
Operational Controls	24
Assignment 3: Router and Firewall Policies	26
General Security Stance	26
Border Router Security Policy	26
WAN Interface Ingress	26

VPN Interface Ingress	28
Firewall Interface Ingress	28
Primary Firewall Security Policy	29
VPN Traffic	30
Border Router Traffic	30
DMZ Traffic	30
Internal Traffic	30
References	33

List of Figures

Figure 1: Network Diagram with IP addresses	21
---	----

List of Tables

<u>Table 1: Customers Group Security Requirements</u>	12
<u>Table 2: Suppliers Group Security Requirements</u>	13
<u>Table 3: Partners Group Security Requirements</u>	13
<u>Table 4: Employees Security Requirements</u>	14
<u>Table 5: Remote Employees Security Requirements</u>	14
<u>Table 6: General Public Group Security Requirements</u>	15
<u>Table 7: General Public Data Flow</u>	15
<u>Table 8: Customer Data Flow</u>	16
<u>Table 9: Employee Data Flow</u>	17
<u>Table 10: Remote Employees Data Flow</u>	17
<u>Table 11: Authentication and Authorization Data Flow</u>	24
<u>Table 12: Logging Data Flow</u>	25
<u>Table 13: WAN Interface ACL</u>	27
<u>Table 14: VPN Interface ACL</u>	28
<u>Table 15: Firewall Interface ACL</u>	29
<u>Table 16: Packet Filter Rule Base</u>	30

Assignment 1: Future State of Security Technology

"All warfare is based on deception ... Hold out baits to entice the enemy. Feign disorder, and crush him."

Sun Tzu, The Art of War

Abstract

Although network perimeter security is not comparable with warfare in terms of morality and the seriousness of the risk to human life, the blackhat threats facing network perimeter security use similar tactics to those used in warfare. Network administrators have several tools in their arsenal for thwarting such attacks such as firewalls and intrusion detection systems. A relatively recent concept developed to compliment existing network defense tools is the honeypot.

Honeypots should not regularly receive any interaction, but their value comes when an attacker happens to interact with it. There are many different forms of honeypots, and each form comes with its own benefits and drawbacks. The benefits are usually in the form of aiding in intrusion prevention, intrusion detection, incident response or research. Drawbacks include resource costs and legal risks. In order to produce results, honeypots must interact with attackers. Honeypots typically accomplish this through deceiving attackers by making them think that the honeypot system would be of value to them. Additionally, some forms of honeypots will take preemptive action against attacks before the compromise of any valuable resources. As has been shown throughout history in the face of similar threats, preemptive and deceptive policies are ethical practices in thwarting the attacks of such a threat.

Even ethical practices are subject to the law; likewise, honeypots are subject to scrutiny by state and federal statutes involving issues such as entrapment, privacy and liability. Without legal precedence to govern the applicability of these issues to honeypots, it is impossible to predict their applicability; therefore, any honeypot deployment does incur legal risk to the honeypot owner.

The purpose of this paper is to analyze the ethical and legal implications of honeypots as well as the technological impact that honeypots will have on the information security community as a whole as well as the impact on those who will have to administer honeypots. In the following sections, honeypots are defined and some different forms of honeypots are described. Next, historical evaluations of preemptive and deceptive policies show honeypots to be an ethical practice in network security. Then, the author's opinions on how entrapment, privacy and liability apply (or do not apply) are discussed. Finally,

the benefits of honeypots are weighed against the drawbacks of honeypots in the context of their use in modern society to show how the various forms of honeypots will make positive impacts in appropriate niches of network security.

Background / Introduction

Defining Honeypots

Defining honeypots is a tricky business. Honeypots cannot be defined by the problem they mitigate since they “do not solve a specific problem” (Spitzner). Instead, honeypots are flexible tools that used for a wide variety of purposes.

Several definitions of honeypots may be found, but for the purpose of this paper the very general definition given by Lance Spitzner is favored since it is a flexible definition applicable to the many forms of honeypots. Spitzner defines a honeypot as “an information system resource whose value lies in unauthorized or illicit use of that resource.” This is a very broad definition and deserves some analysis. First, an information system resource can represent a variety of things. The most obvious resource would be a computer connected to the network. Other hardware components such as routers and switches, along with data such as user names and passwords, are also resources in an information system. The next part of the definition describes how the resource becomes valuable. The resource is valuable only when used in an unauthorized manner.

As the general definition implies, honeypots are a broad category of security tools. This broad category can be further divided into the subcategories of low-interaction honeypots and high-interaction honeypots.

Low-interaction Honeypots

This category of honeypots provides limited amounts of interaction for attackers through emulation of a service or of an operating system. For example, a low-interaction honeypot may emulate a Bind DNS server by actually performing some name resolution services or it may just provide fake responses to queries.

Because low-interaction honeypots provide emulation of services or systems rather than providing actual services or systems, this category of honeypots provides a barrier between an attacker’s action on the honeypot and the rest of they system. Additionally, for many low-interaction honeypots, this emulation layer provides the flexibility of emulating several types of services or systems. For example, the [Honeyd](#) low-interaction honeypot can be configured to create multiple virtual hosts on a network running specified services and emulating specified operating systems.

High-interaction Honeypots

In contrast to emulating services and systems, high-interaction honeypots provide real operating systems and services to attackers. Here, to deploy a Bind

DNS server honeypot, a Linux system needs to be setup with a Bind package installed and configured. Since no software is written to interact with attackers, no assumptions are made of the attackers' activities. Providing real systems for the attacker to interact with provides no limits outside of what the attacker would normally expect for the system they think they have compromised.

Value of Honeypots

The value of honeypots exists in the unauthorized use of the honeypot resources. When properly configured, honeypots only receive interaction during unauthorized activity. Due to the hopefully infrequent nature of this activity, data sets produced through honeypot logs should be relatively small. Small data sets can be a weakness as well as strength. Honeypots will not catch all unauthorized activity, but the activity it does receive is very significant. The logs provided by honeypots can be useful for researching hacking techniques, intrusion detection and incident response.

Honeypots are useful to the research community since they can provide a safe zone to study hacker activity without risking valuable information resources. Additionally, since the logs are very concentrated with hacker activity, the researchers do not have to filter through extra data to study the hacking techniques taking place.

Traditionally, the intrusion detection problem has been confined to IDS-s, most of which have a higher than desired false positive rate and require consistent updates. When used as IDS-s, honeypots should provide very low false positive rates. Any activity directed at the honeypot can be considered as an early warning sign of unauthorized activity. The logs can be analyzed to either track down where the attack is coming from or to try and determine the intent and nature of the attack. Honeypots are not meant to replace IDS-s. They provide an excellent additional layer of protection related to intrusion detection, but since they only track traffic directed specifically at the honeypot itself, they will not catch attacks directed at other resources in the network.

A common problem with incident response is filtering through all the audit log information in order to determine what exactly happened to the system. If the attacker happened to attack the honeypot first before moving on to the production system, the honeypot logs may provide information on the nature or goal of the attack. This will provide some correlation with the logs on the production system. An additional benefit of honeypots in general arises from the fact that in most settings (except possibly research), the honeypot resources are not mission critical and can be taken offline to be studied. Additionally, the limited amount of interaction with honeypots means that their configuration state is consistent. These two aspects of honeypots make them excellent resources for incident response investigations. The honeypot resources can be analyzed to see any changes that were made to the configuration and the logs can be analyzed without much filtering to determine the nature and/or source of the

attack. All of this analysis can be done offline with minimal impact on the mission of the organization running the honeypot. The honeypot may also provide good evidence for legal action as long as the evidence was collected in a legal manner. The legal aspect will be explored later in this paper.

In addition to the above-mentioned general benefits of honeypots, low-interaction and high-interaction each have their own advantages. Low-interaction honeypots are usually easier to deploy (usually just involving software installation with a bit of configuration). The emulation also limits what actions an attacker can take on a honeypot, thus reducing risk. High-interaction honeypots have the advantage of making no assumptions about an attacker's activities and can thus provide more information about attacks. However, high interaction honeypots incur more risk.

Tarpits

Tarpits are a special type of honeypot that deserve extra attention. Tarpits are a specialized low-interaction honeypot designed with the purpose of slowing down or stopping attacks.

Background

The original concept of tarpits started in response to the Code Red worm (Liston). The Code Red worm would scan for web servers listening on port 80. All of this scanning took up a lot of bandwidth, so Tom Liston came up with a hypothetical program that would slow down this scanning activity. The idea became a proof of concept and has since turned into the Labrea Tarpit project, named after the Labrea Tarpits in California.

How Labrea Works

The Labrea Tarpits in California are one of the world's most famous fossil areas. Animals would become trapped and preserved in the warm sticky asphalt (Page Museum). Similarly, Liston's Labrea Tarpit is designed to trap network worms by becoming sticky. It becomes sticky by taking advantage of the flow control mechanism built into TCP. For this reason, tarpits are referred to as sticky honeypots.

Tarpits are setup to monitor a network's unused IP address space. The tarpit assumes that anything directed at an unused IP address is malicious. The Labrea Tarpit does this by listening for unanswered ARP requests. When the tarpit sees that nobody is responding, it then forms its own response with a made-up MAC address that does not exist on the local network. From then on, the router forwards packets destined for the unused IP address to the tarpit.

The tarpit listens for, and acknowledges, TCP SYN packets. After the acknowledgement, the attacker may send the last acknowledgment in the three-way handshake along with some data. At this point, the tarpit may respond in several ways. One is to not respond. In this case, the sender will eventually

timeout and may try to send the data again. Usually, after timing out once, the sender will increase the timeout length by an amount determined by the TCP stack implementation. This increased amount may be a constant value (add 6 seconds every time), or it may increase exponentially (6, 12, 24 ...). Another way in which the tarpit could respond is to respond with a small window size (perhaps zero). A window size of zero is normally meant to tell the sender that the buffer is full and that they should wait before sending any more data. Under normal TCP circumstances, the receiver is supposed to advertise when the window becomes greater than zero. Due to the reliable nature of TCP, the sender will send a window probe after a certain period to ensure that the window advertisement is transmitted reliably. The recommendation for the timeout period is two minutes (Postel). The tarpit can then respond to each of these window probes again with a window size of zero. Unless the sending TCP stack is designed to detect tarpits, this interaction should be indefinite without manual interdiction, something that an automated worm will not have.

Value

The value of tarpits can be linked to many of the general benefits of honeypots. Tarpits can be used as an additional form of intrusion detection through alerting administrators when trapping an attacker. They can also serve as decoy against reconnaissance TCP scans and deterrence to automated worms. Tarpits also enjoy most of the general benefits of low-interaction honeypots. Tarpits actually slow down the interaction of the attacker with the honeypot system, thus not allowing the attacker to perform any actions on the honeypot system. The Tarpits are limited to TCP since they rely on flow control; therefore, any attacks over protocols other than TCP are not affected by tarpits.

The HoneyNet Project

Another type of honeypot whose gaining prominence in information security research requires a deeper look is the HoneyNet Project. HoneyNets are a specialized type of high-interaction honeypot designed with the purpose of providing an environment for monitoring and researching attacker activity.

How HoneyNets Work

HoneyNets are purely an architectural concept. HoneyNets are simply networks with honeypots on them. Providing a highly controlled environment for blackhat activity is central to the honeyNet concept, thus honeyNet architectures have two critical requirements for establishing this controlled environment. The first and utmost requirement listed in the HoneyNet Project's definitions paper is data control. With the goal of honeyNets being the research of blackhat activity, it is actually desirable to see honeypots on a honeyNet become compromised. If not properly controlled though, this compromised honeypot may become a liability if future blackhat activities on the compromised box are attributed to the owners of the honeyNet. The trick to this is to control the environment without tipping off the attacker that they are in a honeyNet. The other critical requirement for establishing a highly controlled environment listed in the HoneyNet Project's

definitions paper is data capture. This requirement is a lower priority than data control, but is essential to the goal of honeynets. To implement this requirement, everything that comes into or out of the honeypot must be monitored without notifying the attacker that they are being watched.

With the development of the Honeywall project, data control and capture are usually implemented using second-generation honeynets as described in the Honeynet Project's GenII Honeynets paper. The Honeywall acts as an invisible bridge between the honeypots and the rest of the network. The Honeywall adds another layer of control on hacker activities since any activities that happen on the honeypots have to flow through the Honeywall. The Honeywall includes network security tools such as a bridging firewall, Snort IDS and snort-inline IPS.

The bridging firewall contributes to the data control goal by limiting the amount of data that can pass from any single honeypot per some reasonable unit of time. This serves to prevent any denial of service attacks originating from the honeynet. The other technology utilized by the honeywall CD is the inline-snort IPS. This system uses Snort IDS to recognize attack signatures along with modified Snort rules to modify the attack packets passing through the Honeywall. This thwarts the attack without tipping off the attacker as to why the attack failed.

Value

Like tarpits, a lot of the value of honeynets is common to all honeypots. Unlike tarpits, honeynets are high-interaction honeypots and thus provide attackers the freedom to act upon real systems. This flexibility provides an ideal environment for researching attacker activities but also incurs more risk since attacker activities are not completely contained in an emulated environment. Honeynets also require more work to implement the data control and data capture goals at the heart of honeynet standards. The Honeywall CD attempts to ease the deployment of second-generation honeynets, but they still require more effort to manage than most network security solutions.

Ethics of Honeypot Use

All forms of honeypots find their value by deceiving attackers and/or proactively defending a network. Deception and preemptively taking actions against attacks before any laws have been broken places these tools into a moral and legal gray area. The purpose of this section is to address the ethical issues surrounding the preemptive and deceptive nature of honeypots.

The definition of ethics given by the Merriam-Webster Dictionary is "a discipline dealing with good and evil and moral duty." The terms "good and evil" along with "moral duty" suggest religious and philosophical undertones. Traditionally, ethics were defined by religious institutions and businesses tended to follow suit with ethical business practices that reflected the social moral value system.

One example from United States history is the Puritan work ethic. Although religion still has a strong role in defining the moral values that are to be reflected in ethical practices, a spreading postmodernist philosophy where truth and moral values are created rather than absolute has moved the role of defining moral values away from religious institutions. In the name of tolerance, religions receive criticism for their exclusive claims to what is right and wrong. Since the source of what defines ethical practices has become less authoritative in modern society's eyes (especially in the Western world), the role of defining ethical practices in society has shifted.

The role has shifted somewhat toward evaluating practices on a historical basis and determining if the practice has had a beneficial impact. One example from world history is the Cold War and the practice of détente versus Raegan's Strategic Defensive Initiative (SDI). The focus of détente was to ease the tensions between the U.S.S.R. and the western world. On the other hand, the focus of SDI was to leverage technological advantages in order to thwart a nuclear attack. Détente's idea of easing the tensions between two nuclear powers to avoid nuclear warfare seemed ideal at the time, but it failed to recognize the expansionist nature of the Communist threat. The expansionist nature took advantage of this détente and considered it as an opportunity for continued expansion. The SDI was a policy shift that transformed the nature of the U.S. defensive policy from reactionary to proactive since the technology would be able to stop launched nuclear weapons before they reached the United States. Although the SDI technology never materialized (there was less need for the program after the economic collapse of the Soviet Union), the shift in policy from reactionary to proactive eventually ended the Cold War. A historical evaluation of SDI reveals that the proactive policy had a beneficial impact on the society as a whole and appears to be an ethical practice in the face of an expansionist threat.

Many parallels can be drawn between the Cold War example and the use of honeypots in defending networks. Blackhat attacks are expansionistic in their own form. Worms try to copy themselves from one computer to another. Attackers may only want to take control of a system to launch future attacks from the compromised system as an additional level of non-attribution. These examples show that the philosophy that a system is not at risk since it is not a high profile system is a philosophy that fails to recognize this expansionist nature of blackhat activity. Additionally, it fails to adapt to the nature of businesses striving to do well and thus becoming higher profile businesses. Instead, as history has shown, a proactive defense posture is an attractive and ethical solution to defending a network perimeter.

When discussing the deceptive nature of honeypots, it is important to keep in mind the target of the deception. Honeypots are designed exclusively to deceive blackhat activities including automated attacks from worms and attacks from human hackers. These circumstances represent threats whose tactics

resemble that of warfare more than normal circumstances where deception is generally unethical. The well-known work by Sun Tzu, The Art of War, addresses the ethics of warfare specifically citing deception as being core to all warfare. Would it be wrong to deceive those engaging in warfare tactics for the sole purpose of protection? It is the opinion of the author that the use of deceptive techniques against attackers for protective purposes is an ethical practice. Historical evaluations of examples provided by honeynets and tarpits that meet the criteria of ethical deceptive techniques reveal a net beneficial impact on society.

An additional ethical issue not covered by any of the previous analysis relates to the liability of honeypots. The liability issue while a legal issue can also be an ethical one. The liability issue applies to high-interaction honeypots where attackers have the freedom to use the honeypot to attack other systems. High-interaction honeypots work well for research, so the liability issue seems to boil down to a responsible research practices issue. In short, researchers should exercise due diligence in their research by utilizing existing security technologies to protect other networks from the research honeypot activities.

Legality of Honeypots

The previous sections have already established the fact that honeypots have a deceptive nature. The first legal issue that comes to mind is entrapment. The other legal issues addressed in this section are privacy and liability. Before diving into the legal issues of honeypots, it is important to note that honeypots are a relatively new technology with no legal precedence. In addition, legal issues may differ from state to state. Honeypot owners need to be aware of any state-specific legal issues that affect honeypot usage, specifically legal issues regarding privacy, entrapment, consent and anything defining a communication device as being illegal. Finally, the author is not a legal expert and this paper will only be an expression of opinions on how the law applies to honeypots.

Legal Definitions Dot Com defines entrapment as a situation “where a police officer or other law enforcement officer induces a person to commit a crime that the person wouldn’t have committed otherwise for the purpose of bringing a criminal prosecution against that person.” Entrapment only applies to the law enforcement community. In cases where law enforcement personnel press charges against a blackhat where a honeypot was involved, the prosecution would have to prove that the honeypot did not induce the person to commit the crime. This should be straightforward since the attacker had the choice whether to attack the system or not. Even if the honeypot was advertising itself to be a vulnerable system to an attacker, saying this is entrapment would mean that blackhats attacking any government systems with weaknesses would have entrapment defenses.

Like any network-auditing tool, privacy may be an issue. Privacy is a very complex and highly debated issue and thus will not be completely explored in

this paper. The main factor determining the effects honeypots have on privacy is how the honeypot is being used. If the honeypot is being used for the purposes of protecting a network, the honeypot falls under the service provider protection exemption in the Federal Wiretap Act. Most uses of production honeypots fall under this exemption. One use of honeypots that may not fall under this exemption is research. The only privacy a honeypot could invade would be the privacy of an attacker illegally accessing the honeypot system. The attacker would not have a legal right to privacy, but the Federal Wiretap Act may make it illegal for a researcher to record the communications of the attacker interacting with the honeypot system. The most common way of addressing this issue is to provide banners where users of the system consent to monitoring. This is not perfect since not all ports may be bannered. The attacker may not understand the language the banner is written in, and different states have different laws dealing with consent to monitoring. Despite these imperfections, banners shows due diligence on the part of the honeypot owner and reduce the risk of being prosecuted for privacy violations.

The last legal issue that honeypot owners need to be aware of is liability. As was mentioned earlier, only high-interaction honeypots have liability risks. Low-interaction honeypots limit hacker activities through emulated systems. The no-limit nature of high-interaction honeypots may allow blackhat to attack other systems from a compromised honeypot. There have not been any legal cases setting precedence where the owner of an insecure system was found liable for damages done against another system using the insecure system. Do system owner's have a responsibility to reasonably secure a system? A similar example might be the owners of a swimming pool. If harm comes to a child in a swimming pool that does not have proper security measures, the owners of the swimming pool may be liable for the harm to that child (Haggard, HINTON V. 2331

ADAMS STREET CORP). The key point here is that due diligence was not performed in securing the resource that has the potential to cause harm. It seems that as long as due diligence is taken in securing the resource, the owners should not be held liable for any harm that resulted from actions taken to bypass security measures that were taken to secure the resource. No system is 100% secure, therefore if owners of systems that have shown due diligence in securing their system are found liable for attacks using their systems, then any system can be found liable in an attack. Due diligence in securing a system should include adequate logging and audit log reviews in case a compromised system was used as a final hop point in an attack. In such a case, the target of the attack will see that attack as originating from the last hop point. Adequate logs will be needed in order to show the attack did not originate from the system.

Impact on Perimeter Security

Questions of legality and ethics are of no importance if honeypots are never used, and they will only be used if they have a beneficial impact. Honeypots

have already had some positive impacts such as discovering new exploits in the wild and holding some worm threads in persistent fake connections. These successes have helped to establish honeypots in the information security community, and their use will only grow with future successes. It is still early to judge the overall impact honeypots will have on the information security industry. The purpose of this section is to discuss how honeypots will affect the information security industry and the system administrators that will have to manage them.

Chen, Gao, and Kwiat attempted to model the impact of sticky honeypots against worms such as Code Red. The model shows that the Labrea Tarpit would have a significant impact in thwarting worm propagation if tarpits were listening on at least 2^{18} IP addresses. The problem with this model is that it assumes that future worms will not account for tarpits. The current trend in writing worms shown by Code Red, Slammer, Sobig and MSBlaster is to make worms small and fast with the goal of spreading as quickly as possible. If this trend continues and tarpits become more prominent, newer worms will be written to account for tarpits to avoid them and continue to spread rapidly. They could do this by implementing their own TCP/IP stacks to look for tarpit signatures (small window sizes), but this would increase the size of worms making them bulkier and slower when copying themselves over the network. Other ways may include some form of distributed scanning before releasing a final version of the worm onto the internet, using results of the scans to preprogram threads of the worm to avoid tarpits. In any case, an increased use of tarpits will force worm writers to start accounting for them and when they do, tarpits will become obsolete against the new smarter worms. In short, sticky honeypots appear to be a beneficial short-term tool in fighting worms, but their reliance on TCP flow control and resulting signature will most likely cause them to become obsolete before enough sticky honeypots are deployed to have a significant impact against worm propagation.

Although honeypots may not have a huge impact on intrusion prevention, their value in detection and response will play a large role in the overall impact honeypots will have on the information security industry. They are not designed to be replacements for typical network and host based IDS-s, but adding a honeypot as an additional layer of security in a production environment gives network administrators an additional tool in identifying false-positives that have traditionally plagued IDS-s. Additionally, activity on honeypots can aid network administrators in identifying false-negatives missed by IDS-s due to inability to download the latest signatures. This additional level of security acts as an enabling tool allowing network administrators to run the network with a stronger security posture in a more efficient manner. In the event that a false negative is not caught in time and a security incident does occur, the relatively small data set provided by the honeypot logging mechanism will be of great value in efficiently determining the nature and scope of the attack.

Honeypots may find their greatest impact in information security research. This appears to be the case since this is where honeypots are already becoming the most prevalent, especially with the growing popularity of the HoneyNet Project. The nature of studying blackhat activity is a very popular topic and appeals greatly to the academic community. Additionally, the value of the research has become more and more valuable to companies that develop intrusion detection, virus scanning and firewall tools. HoneyNets are ideal tools for this kind of research whose use will increase along with the increase in demand of this research topic.

Conclusion

Honeypots are systems that find their value in their unauthorized use. This value applies to the preventive, detection and response phases of information security. Their benefit in providing relatively small data sets highly concentrated with logs of unauthorized use is also their weakness in that they only record a limited view of data. Due to their limited view, honeypots are not replacements for existing perimeter security technologies, but instead should be used as an additional layer of security that will aid administrators in providing stronger security in a more efficient manner. Although honeypots are preemptive and deceptive in nature, historical evaluations of preemptive and deceptive policies in confronting threats similar in nature to blackhat activities have proven that honeypots can be ethical tools when used for the benefit of society as a whole. Without legal precedence, implementing honeypots does incur some risks of legal action which should be considered in the risk management process. The smaller risk of low-interaction honeypots make them ideal tools to use in production environments for intrusion detection and incident response. High-interaction honeypots such as the HoneyNet Project require more resources for administration and also incur more risk. Thus, they probably will be confined to the research community. Although the long term benefits of tarpits will most likely diminish as automated worm attacks become smarter, the short term benefits of slowing down traditional worm propagation should not be overlooked. Organizations deploying tarpits will contribute toward the goal of seriously affecting the propagation rates of the worm threat thus showing their commitment to the information security community.

Assignment 2: Security Architecture

The purpose of assignment two of this practical paper is to propose a network security architecture for a small company whose mission is to market fortune cookie sayings worldwide. The company's name is GIAC Enterprises. They employ fifty people distributed among a headquarter office and four international branch offices. All of GIAC's sales are done via the internet.

Access and Security Requirements

Several different roles of interaction with the GIAC enterprises network exist with different access requirements. These roles include customers, suppliers, partners, employees and the general public. Employees are composed of those located at the headquarters along with remote users at branch offices and will be treated as two separate groups when defining their access requirements. Group descriptions and their access requirements are enumerated in the following subsections. Additionally, the security services required for the interface for each group and the GIAC network are described in order to lay out the foundation for the security architecture of the network design.

Customers

The customers group interacting with the GIAC enterprises network is composed of companies or individuals that purchase fortune cookie sayings from GIAC via the Internet. Customers make purchases through GIAC website. They only need direct access to the website. The table below describes the security services and applications required for the interface between customers and the GIAC network.

Table 1: Customers Group Security Requirements

Security Service	Data	Description
Confidentiality	Sensitive Customer Information	Ensure information such as credit cards and banking information are not disclosed
	Fortune Cookie Sayings	Ensure only the buyer can see the bulk order of fortune cookie sayings
Integrity	Sensitive Customer Information	Ensure customer information is processed exactly as provided by the customer
	Fortune Cookie Sayings	Ensure customer's receive an unaltered product
Authentication	Authentication of Webserver	Give customer's some assurance they are interacting with GIAC
Non-repudiation	Contract	Provide some assurance that customer cannot deny making a purchase
Availability	Customer Interface	GIAC's profits come from purchases made through the internet, thus the availability of this interface (webserver, backend) is vital.

Suppliers

GIAC enterprises obtain fortune cookie sayings from various suppliers in electronic format. Management typically makes purchases through the supplier

website. In addition, there is a supplier web interface in case the supplier does not have the capability to protect GIAC sensitive purchase information through their web interface or if the supplier does not have a web interface at all. For this case, the supplier will need access to the supplier web interface. No fortune cookie sayings are directly added to the GIAC database by the suppliers themselves. Instead, all supplied sayings are placed in an intermediate database where they are scrutinized to ensure that the supplier did not include duplicates and to ensure that the sayings do not have any foul language in them. After analysis, the sayings are then put into the GIAC database by GIAC employees.

Table 2: Suppliers Group Security Requirements

Security Service	Data	Description
Confidentiality	Fortune cookie sayings	Ensure that only GIAC will be able to see the supplied fortune cookie sayings.
Integrity	Fortune cookie sayings	Provide functionality to aid in integrity preservation for transactions from suppliers.
Authentication	Transactions	Ensure that the transactions originate from who they say they originate from.
Non-repudiation	Delivered product	Provide assurance that suppliers cannot deny the state of any delivered products.

Availability is not as big a concern since supplier transactions are made much less frequently than customer transactions due to the increased size of supplier transactions and the fact that sayings can be reused multiple times for different customers.

Partners

The international businesses in the business partners group translate GIAC fortune cookie sayings into other languages and resell them. Since GIAC engages in international business-to-business activity, the GIAC legal team will help to ensure that e-business practices (including security) recommended by the International Chamber of Commerce (ICC) are adhered to. As part of the business partnership agreement, partners are granted full access to all fortune cookie sayings in the GIAC database. Instead of providing direct access to the main GIAC database, partners are granted access to a partner SSL enabled website, which interfaces with the GIAC database. For this SSL connection, the web server in addition to requiring a password on the web interface will request a client certificate.

Table 3: Partners Group Security Requirements

Security Service	Data	Description
Confidentiality	Fortune Cookie Sayings	Only business partners will see the fortune cookie sayings during transactions.
Integrity	Fortune Cookie Sayings	Fortune cookie sayings are unaltered throughout a transaction.
Authentication	Business Partner Identification	Partners must authenticate themselves in order to gain access to partner database

Availability	Business Partner Interface	As per the partnership agreement, the availability of the GIAC database through the partner database is a high priority.
Non-repudiation	Business Partner Transactions	As per the partnership agreement, business partner transactions must be logged in a reliable manner. Also ensures that partners cannot deny making transactions that may appear concerning.

Employees

The employees group is composed of users on the internal network at the headquarters office. All employees have internet access and email services. The internet access is provided through a proxy server. Additionally, the principle of least privilege is enforced. Users are divided into groups and access to some resources on the internal network is based upon group based access controls. Sharing between groups can be requested through the system administrators setting up shared folders on the file server or granting access to individuals for certain application services.

Some employees (system administrators) will require SSH access to the firewall and router for administrative purposes.

Table 4: Employees Security Requirements

Security Service	Data	Description
Access Control / Authorization	Workgroup resources	Enforcing the principle of least privilege, access to workgroup resources should only be granted to members of that workgroup except when resources are labeled as shared
Integrity	Critical Resources	Integrity controls will be used to protect user created resources that are labeled as being critical to GIAC
	Security Resources	Integrity controls will be used on appropriate security related resources to add assurance that the network is in a good state and to aid in handling potential security related problems
Authentication	GIAC Resources	Employees must authenticate themselves in order to access any resources on the GIAC network
Confidentiality	SSH Sessions	Configuration sessions for system administration should be encrypted so that only system administrators can access the routers and VPN gateways
Availability	Application Servers	To ensure minimal production lost, availability of application services is important

Sales/Teleworkers

The sales/teleworkers (remote employees) group is composed of users at remote branch offices as well as employees on travel. Users in this group have access to application services as well as the internet and email. To gain access to services located at the corporate headquarters, remote users must connect to the headquarters LAN through a VPN. To enforce principle of least privilege, remote users are assigned to groups so that they are only granted access to services available to their respective group.

Table 5: Remote Employees Security Requirements

Security Service	Data	Description
------------------	------	-------------

Access Control / Authorization	Workgroup resources	Enforcing the principle of least privilege, access to workgroup resources should only be granted to members of that workgroup except when resources are labeled as shared.
Authentication	Headquarters LAN access	Remote users must be authenticated before being allowed access into the corporate LAN.
Integrity	Communications	Remote user communications shall not be altered without notification to the users.
Confidentiality	Communications	All corporate communications over public telecommunications networks shall be encrypted.

General-Public

The general-public only has access to the corporate website handled through the public relations department and can email GIAC employees. They do have access to supplier and customer interfaces through the web server, but if they engage in any transactions, they are not acting as a member of the general public group rather as a member of the suppliers or customers group.

Table 6: General Public Group Security Requirements

Security Service	Data	Description
Availability	Web Page	For marketing purposes, it is important to keep the web page up and running
Integrity	Web Page	It is important for the company's image that any web page defacement be prohibited

Data Flows

Now that the access requirements have been set, the data flows can be enumerated to aid in transforming the access requirements into a functional network architecture. Each group is addressed with a table enumerating the services the group will be accessing along with associated ports, protocols and descriptions of the protocols about how they fulfill access requirements of the required protocols. All data flow entries are assumed to represent data flow in both directions unless explicitly stated.

General Public

The general public accessing the web site may make requests to the DNS server in the DMZ for the web server address. The IP address for the HTTP proxy is assigned to the 'www' name, thus the DNS server will respond with the IP address of the HTTP proxy. The following HTTP requests will then be directed at the HTTP proxy server, which will perform any necessary filtering and then forwarding the request to the main HTTP server on the internal network. The main HTTP server will then service the request by sending the requested data to the proxy server, which will cache the data and forward it on to the requesting external host.

The general public sending email to GIAC employees may go through a similar

DNS resolution process for the mail relay. Mail is sent to the SMTP relay on the DMZ, validated by the relay, and then sent to the main SMTP server on the internal network. Outgoing mail also travels through the SMTP relay before going out to general public mail servers.

Table 7: General Public Data Flow

Source	Destination	Port(s)/Protocol	Description
General Public	Reverse Proxy HTTP Server	80/TCP (HTTP)	Anybody can connect to the public website.
General Public	SMTP Relay	25/TCP (SMTP)	Anybody can send email to internal employees.
SMTP Relay	Internal SMTP Server	25/TCP (SMTP)	Mail is forwarded from the relay to the main SMTP server.
General Public	External DNS Server	53/UDP (DNS)	Allows external users to query the address of the webserver. 53/TCP is not needed since responses should be small and not require large responses. Additionally, all externally requested zone transfers (through 53/TCP) will be blocked.

Customer

Customers have the same data flows as the general public as well as a secure data flow channel for servicing orders provided by SSL. SSL traffic is forwarded by the proxy server to the web server. The SSL related data flows available to customers are described in the table below. Additionally, the internal web server accesses the GIAC database on the internal network in order to service customer transaction requests.

Table 8: Customer Data Flow

Source	Destination	Port(s)/Protocol	Description
Customer	Reverse Proxy HTTPS Server	443/TCP (HTTPS/SSL)	Customers can make purchases through a web interface protected by SSL.
Reverse Proxy HTTPS Server	Internal HTTPS Server	443/TCP (HTTPS/SSL)	SSL traffic is forwarded by the proxy server to the main internal web server.
Internal HTTPS Server	GIAC Fortune Cookie & Accounting Databases	3306/TCP,UDP (MySQL)	For retrieving customer orders.

Supplier

Management making purchase orders from a supplier's website and thus downloading orders through SSL performs most supplier transactions. This data flow is represented by the employee group's data flow section. In cases where the supplier performs a supply transaction through the GIAC website, the supplier will push the orders through a secure web interface protected by SSL. Supply transactions made in this manner are routed through an application server on the internal network for processing. The supplier group has the exact same set of data flows as the customer group above.

Partner

Partners access the GIAC database through a partner website interface on the

GIAC web server through SSL. After authenticating, partners can download fortune cookie sayings through the website interface, which accesses the main GIAC database to retrieve fortune cookie sayings. The partner data flows are exactly the same as the supplier and customer data flows.

Employees

All employees have a data flow that provides both internet and email access through proxy servers. Additionally, all employees are allowed to send ICMP echo request packets to the network to see if hosts are alive on the Internet. To make DNS requests, employee host machines are configured to direct requests to the DNS server in the GIAC DMZ. This server will then resolve requests through additional requests with external DNS servers or through pulling data from a cache. With resolved IP addresses, employee HTTP requests are directed through the HTTP proxy server where requests will be filtered and then passed on to external web servers upon passing the filtering mechanism. Responses to the HTTP requests come through the HTTP proxy server, filtered again and then finally passed back to the employee hosts. For email, employees will access the main mail server on the internal network. The main mail server will utilize an SMTP relay on the DMZ to service both inbound and outbound email requests. The rest of the employees data flows within the network will be described in the defense-in-depth section.

System administrators will need SSH to administer the border router and VPN gateways while the database administrators need SSH and SQL client access to administer the database servers.

Table 9: Employee Data Flow

Source	Destination	Port(s)/Protocol	Description
Employees	Internet	ICMP Echo Request	Used to check if hosts are alive on the Internet
Employees	DMZ DNS Server	53/UDP (DNS)	Used to resolve external DNS names
DMZ DNS Server	Internet	53/UDP (DNS)	Requests from the DMZ DNS Server to outside DNS servers (only done for local net)
Employees	Proxy server	8080/TCP (HTTP)	Allows internal users to access resources on the internet through the proxy server listening on 8080
Proxy Server	Internet	80/TCP (HTTP) 443/TCP (HTTPS/SSL)	Forwards filtered HTTP(S) requests to the internet
System Administrators	Border Router, VPN Gateway	22/TCP (SSH)	System administration activities
System Administrators	Internet	ICMP Echo	System administrator may ping from the routable part of the networks
Database Administrators	Databases	22/TCP (SSH)	Configuring the databases

Database Administrators	Databases	3306/TCP,UDP (MySQL) 22/TCP (SSH)	Configuring the databases through an SQL client or SSH.
Employees	Internal Mail Server	25/TCP (SMTP)	Allows internal users access to email services
Internal Mail Server	SMTP Relay	25/TCP (SMTP)	SMTP is forwarded to the mail relay before leaving the network (headers are stripped)
SMTP Relay	Internet	25/TCP (SMTP)	Outgoing external mail

Sales/Teleworkers

The sales and teleworkers group accesses the shared resources on the internal GIAC network through a VPN tunnel. The rest of the data flow for sales and teleworkers will be addressed by the defense-in-depth section.

Table 10: Remote Employees Data Flow

Source	Destination	Port(s)/Protocol	Description
Sales/Teleworkers	VPN Gateway	500/UDP (IKE)	Key negotiation for VPN establishment
Sales/Teleworkers	VPN Gateway	IP 50 (ESP)	VPN traffic

Architecture Components

Using the enumerated data flows and access control requirements, this section describes the selected critical components necessary to fulfill both the functional and security requirements of the GIAC network. Additionally, the following list includes some principles used in making network design decisions.

- Whenever feasible, show a commitment to the information security community through being a good net neighbor and tarpitting worms.
- The security architecture of the network should be flexible such that a compromise of one aspect of the architecture does not result in the compromise of mission critical resources.
- Use existing standards and protocols over custom solutions.
- Whenever feasible, use open source solutions.
- Encrypt any information passing through public lines.
- Enforce the principle of least privilege wherever possible.
- Prevent outside users from mapping the internal network.
- The existing GIAC IT staff has experience with Cisco IOS, Linux and BSD. OpenBSD is preferred over other Linux/BSD solutions due to the project's proactive security stance and support for cryptographic hardware products.

Filtering Router

The Cisco 2691 Multiservice Platform was chosen to serve as the border router for the GIAC network. The core issues considered in making this decision are in the following list.

- Router must service the T1 WAN connection to the ISP. The router should be upgradeable in case of a WAN link upgrade.

- Router must be able to provide static filtering services as the first line of defense for the GIAC network perimeter.
- The router should be able to handle both the filtering and connection services without a noticeable performance hit.
- The GIAC IT staff is familiar with Cisco IOS.

The Cisco 2691 has a modular design capable of handling a variety of WAN connections. If GIAC upgrades the WAN connection from T1 to something else, only the WAN module on the 2691 needs to be upgraded rather than the entire platform. The 2691 also uses Cisco IOS, which will provide a recognizable configuration interface to the existing GIAC IT staff and will allow them to filter incoming and outgoing traffic. Finally, the processing power of the 2691 should suffice in meeting the performance requirements. Cryptographic processing involved with VPNs is offloaded onto a dedicated VPN server.

The router will have two interfaces into the GIAC network. One of these interfaces will be for VPN traffic, while the other interface is for all other traffic. Only ESP traffic will be allowed to flow in or out of the VPN interface. The other interface will only allow the services required by the access requirements in the above section into the internal network. No stateful inspection will be performed at the filtering router.

Firewalls

The GIAC network utilizes Packet Filter (PF) running on OpenBSD 3.6 configured for stateful inspection and packet normalization as the next line of defense behind the filtering router. The OpenBSD system will have the following four interfaces: border router interface, VPN interface, DMZ interface, and an internal interface. This firewall will block all traffic not explicitly allowed by the data flows and access control requirements explicitly enumerated in the previous sections or in the defense in depth section.

Two boxes on the DMZ are devoted to running Squid 2.5. The first box runs Squid as an outbound proxy server on OpenBSD 3.6 in non-transparent mode. Requiring all internal hosts to go through the Squid proxy for outbound connections prevents any infected internal hosts from connecting to the internet without explicitly knowing where the proxy server is located. Additionally, content filtering is done at the proxy server helps to enforce the company policy against accessing inappropriate material on the internet. The connectivity through the proxy server is limited to only users on the private GIAC network using Squid's access control features. The other box will run Squid 2.5 in reverse proxy mode (httpd-accelerator) on an OpenBSD box. The reverse proxy ensures that external hosts do not directly access the main web server. The proxy server passes through all SSL traffic.

VPN

An OpenBSD 3.6 box configured to serve as a VPN gateway provides VPN access to all branch offices and teleworkers. All VPNs connecting to this server use IPsec in tunnel mode to connect to the server. Having the VPN separate from the border router and firewall removes the cryptographic processing load from these critical components and moves it onto the dedicated VPN gateway. Additionally, the traffic coming through the VPN gateway must go through the stateful inspection firewall before accessing any resources on the GIAC network.

The VPN tunnels will use a shared secret key to protect the VPN traffic. Shared secret is used here since small number of tunnels required does not warrant the establishment of a public key infrastructure. The preferred algorithm for the tunnels is AES 128, but any AES algorithm, triple DES or Blowfish is acceptable. These options provide excellent security with reasonable performance.

Network based IDS sensor(s)

An OpenBSD 3.6 box running Snort 2.1.2 with multiple inputs from taps throughout the network is used for intrusion detection. A tap is placed at the decrypted side of the VPN network and on the DMZ network. No tap is placed outside the filtering router or between the filtering router and the stateful firewall, which means that not all attacks directed at the GIAC network will be noticed. The already over-burdened IT staff only needs to concern itself with attacks that have made it through to the internal network. The OpenBSD box itself is on the internal network and is only accessible by the system administrators.

Additional Components

Honeypot / Tarpit

An additional box on the DMZ network will run the Honeyd honeypot. The honeypot is configured to simulate a Linux host running HTTP and SMTP services. The domain names assigned to these virtual services are different enough from the real services so that no legitimate user should mistakenly attempt to use the honeypot services instead of the real services. This virtual Linux host will have one of the remaining IP addresses. All other remaining IP addresses are tarpitted.

Intrusion Prevention System

The most critical resource related to the mission of GIAC enterprises is the database system storing the fortune cookie sayings. In order to provide extra protection for this system, all traffic directed at the database servers must go through a Linux box running Snort Inline IPS. All traffic directed at the database server from customers, partners and suppliers should come through the web site and be composed of a well-defined set of queries. Because of this, there should be a low rate of false positives detected on this traffic; therefore, traffic

from these hosts should go uninterrupted by Snort blocking a valid request that happened to register as a false positive.

SMTP Relay

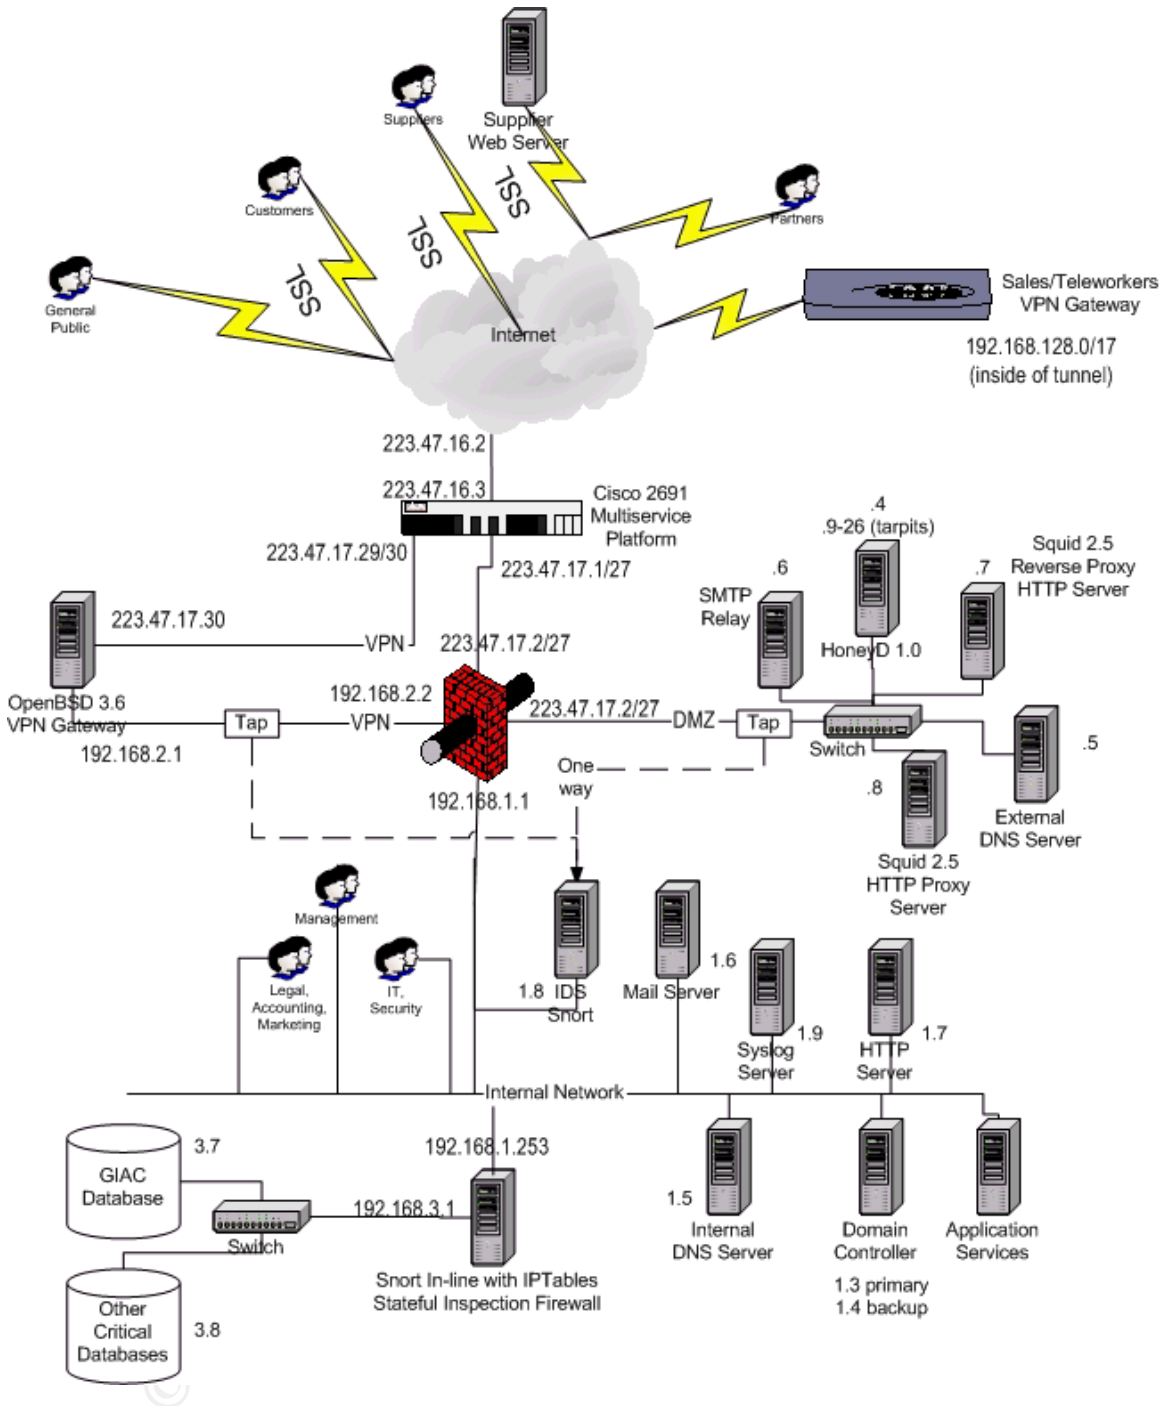
The SMTP relay is configured so that it will not act as an open relay for use by anybody on the internet, instead it will only be able to relay mail for the GIAC network. By utilizing an SMTP relay on the DMZ network, the main mail server is not directly exposed to the internet. In order to compromise the main mail server, an attacker would have to gain access to the DMZ network first and then compromise the mail server from there. The additional hop increases the chances that the attacker will be caught. In the event that the SMTP relay has been compromised, the relay only needs to be patched and the configuration restored since it does not actually handle any of the mail messages. Additionally, the relay strips the mail headers as an added level of protection for the internal mail server.

Network Diagram and IP Addressing Scheme

The GIAC network has been allocated 30 public IP addresses in the CIDR block of 223.47.17.0/27. The network and addressing diagram are in figure one below. The VPN address block of 223.47.17.28/30 falls entirely within the address block on the other interface of the border router, but traffic will route properly to the VPN gateway since CIDR routing uses a “longest match wins” policy when deciding which interface to route traffic (Davie).

Figure 1: Network Diagram with IP addresses

© SANS Institute 2000 - 2005



Implementing Defense in Depth

As was stated earlier as a guiding principle in designing the security architecture of the GIAC network, compromise of one component of the GIAC network security architecture does not directly compromise critical resources. Motivated attackers must work through several layers of security to reach these resources, thus increasing the chance of detection and ultimately thwarting the attack.

Intrusion Detection

One key element in providing defense in depth is intrusion detection. In order to meet the GIAC functional requirements, the network must allow some traffic in and out of the network. This traffic opens up the network to the possibility of attacks coming through the open traffic channels. Multiple one-way taps that feed into the Snort NIDS provide several places where these attacks may be detected.

VPN Segment

The branch offices and teleworkers are probably more susceptible to compromise than the GIAC network; therefore, the VPN connections should not be completely trusted. Malicious traffic coming through the VPN tunnels represents a serious problem since it means that either another GIAC system has been compromised or a GIAC employee is attempting to attack the network using insider knowledge. Placing a NIDS tap on the VPN network rather than behind the firewall contributes to defense-in-depth by analyzing all the traffic coming in on the VPN segment to determine if hosts on the other end of the VPN tunnels are compromised.

DMZ Segment

The core objective for the NIDS to accomplish is to augment the security architecture of the network by protecting service network against external attackers. Placing a tap on the DMZ segment behind the stateful firewall allows the NIDS to analyze filtered traffic coming into the services network from the internet. Due to the placement behind the firewall and filtering router, the filtered traffic will not be analyzed (although the router and/or firewall will log it).

Intrusion Prevention

The most sensitive information in the GIAC network is stored on the critical resources network. Due to the criticality of the information, this network requires an added layer of protection to prevent attacks from compromising the critical resources, even from within the GIAC network. All traffic destined to the critical resources network passes through the snort inline IPS. Attacks directed at the critical resources, which have known signatures in the snort signature database will be thwarted by this IPS. The defense-in-depth added by the IPS here is that it prevents attacks that match known signatures that either come by an attacker managing to force the web server to make a malicious request through SSL or from an already compromised host inside the network (could be compromised by a malicious insider).

Host/Service Hardening

In case an attack does make it to one of the host systems on the network, the chances of the attack actually exploiting the target system on the GIAC network is reduced by security hardening each of the hosts in the network.

Hardened Windows Images

Each of the general user hosts are set up using a “secure” image of a Windows 2000 machine. The images are created from an up to date and patched version of Windows 2000 with automated password rule enforcement enable in the Windows policy. Additionally, the Windows NT Lanman hashes are disabled and unnecessary tools such as telnet are removed from these images.

Bastion Router

The border router is a key element to the security architecture of the GIAC network. To make matter worse, it is directly exposed to the internet. In order to reduce the risk of the border router being compromised, the router is security hardened by disabling unneeded services, using secure passwords, and using SSHv2 instead of telnet for remote configuration.

Services

The GIAC network has many services, both on the DMZ segment and the internal network. Utilizing OpenBSD to run the services adds some assurance that the services are running on a secure platform (including stack execution protection to protect against classic buffer overflows). Any unnecessary services are disabled and strong passwords are enforced. Additionally, all services are run at the lowest privilege level possible. Policy should dictate that system administrators monitor Bugtraq and possibly other underground lists to ensure that services are patched when vulnerabilities are found.

Deception and Deterrence

An additional layer of defense is provided by making an attacker’s job more difficult through the use of a virtual web server and mail server on a honeypot system. An attacker looking for the “low hanging fruit” in the GIAC system may accidentally stumble across the honeypot system and thus spend a lot of time trying to attack a system with little value. These systems are not for production use; therefore, any interaction with these systems is considered to be an early warning sign of a possible attack. Additionally, each of the IP addresses in the pool of unused IP addresses are configured as tarpits to assist in the detection and deterrence of automated attacks from worms. A positive side effect to detecting and deterring automated attacks is that employing tarpits shows the companies dedication to the information security community.

Authentication and Authorization

In order to enforce the principle of least privilege, users must authenticate themselves to prove they are who they say they are and be authorized to ensure that they are allowed to use the resources they are attempting to use. In order to ease this process, Kerberos is utilized to provide single sign-on authentication services while LDAP is utilized to provide authorization services. For

applications that do not interface with these protocols, other authentication mechanisms such as passwords are used at the applications themselves. Before any authentication takes place, all services should provide a banner stating that the service is for authorized use only and that further use of the resource is consent to monitoring. The following data flow table is added to allow remote users to authenticate and authorize themselves to the GIAC network.

Table 11: Authentication and Authorization Data Flow

Source	Destination	Port(s)/Protocol	Description
Remote users via VPN Gateway	Domain Controller	389/TCP (LDAP)	Lightweight Directory Access Protocol
Remote users via VPN Gateway	Domain Controller	88/UDP (krb5)	Kerberos version 5 (ticket requests)
Remote users via VPN Gateway	Domain Controller	749/TCP (kpasswd)	Allows users to remotely change Kerberos passwords (under unix)
Remote users via VPN Gateway	GIAC Internal Network Resources	Other required Traffic	Other application specific traffic may be required so that remote users can access specific application services on the internal network

Anti-Virus

Viruses are always a threat to network, especially when teleworkers may bring their laptops into the GIAC network. To protect against this threat, a company policy requiring scanning of laptops before connecting them to the network should be enforced. The web server, mail relay and file server on the GIAC network are all configured to use ClamAV to prevent virus infections on the network. All windows clients will be configured to use AVG instead. AVG was chosen since ClamWin anti-virus program for windows does not have a graphical interface, which does not typically work well with the general windows users.

Operational Controls

Several operational controls are also in place to provide further defense-in-depth for the GIAC network security architecture. First, all employees go through some security awareness training when starting work with GIAC as well as annual updates to the training. Employees are trained in recognizing security incidents and reporting them to the proper people. This initiates a well-defined incident handling process, which may also be initiated by things such as IDS alarms or service outages. An incident response team including representation from system administrators, legal team and management team handle the incident response process.

Since GIAC enterprises is a relatively small network, system administrators take on the responsibilities of information security officers as well. This includes

some additional security critical responsibilities such as enforcing user addition/removal procedures and auditing procedures. The system will undergo periodic audits to ensure that the network security architecture is keeping up to date with the ever-changing information security field. The audits should roughly follow a combination of the NSA INFOSEC Assessment Methodology and NSA INFOSEC Evaluation Methodology. Syslog reviews are also included in the auditing process. The Syslog server is on the internal network accessible to the system administrators. With this placement, syslog traffic will need to travel through the border firewall to the syslog server.

Table 12: Logging Data Flow

Source	Destination	Port(s)/Protocol	Description
Any GIAC component on the network	Syslog server	514/UDP (syslog)	Audit logs

© SANS Institute 2000 - 2005, Author retains full rights.

Assignment 3: Router and Firewall Policies

General Security Stance

The general security stance of the GIAC network is a low risk stance that makes compromising critical resources difficult for a determined attacker. In general, only traffic that is explicitly required by the data access control requirements listed in the previous assignment is allowed into or out of the network. Locking down this traffic is accomplished by building filtering rules upon a default deny approach. Unless the traffic is explicitly allowed, it is blocked. The remaining required traffic introduces risk into the network, much of which has already been addressed in the defense-in-depth section above.

Border Router Security Policy

The border router is configured to block all traffic not explicitly allowed onto the GIAC network through the use of static packet filters. The border router filtering policy includes ingress and egress filtering rules to lock down the traffic to only what is necessary. The egress filtering rules ensure the good net neighbor policy along with making the job of an attacker communicating outside of the network more difficult in the event that one of the internal hosts do become compromised. Any packets that are blocked will not trigger an admin prohibited ICMP message. Instead, packets are silently dropped. This configuration comes from the “no ip unreachable” configuration option. The router also prevents any source routing packets from entering the GIAC network using the “no ip source-route” command. Additionally, since the IDS is not listening outside of the network, the border router will also log any “odd” traffic to flag any potential threats. DNS, SMTP and Web Server are some of the top UNIX vulnerabilities according to the SANS top twenty list; therefore all DNS, SMTP and web traffic is logged. VPN traffic represents sensitive network traffic; therefore, it is also be logged.

WAN Interface Ingress

First, all noise traffic is filtered. Most traffic that looks as if it could have been spoofed is logged to flag potential attackers. This includes traffic with source addresses allocated to multicast, loop back, local link block, the GIAC allocated IP addresses and any addresses that have not been allocated (IANA). The private network addresses are rejected but not logged since they may have come from a network configured incorrectly.

Next, all VPN traffic destined to the VPN gateway is permitted. Since teleworkers and branch offices do not all have static IP addresses, any valid source address is permitted to connect to the VPN gateway (from the router's point of view).

Next, all traffic that is allowed to initiate connections with services on the DMZ are permitted through. Then, traffic that corresponds to responses from requests originating on the DMZ are allowed through. The largest hole in the access list is in this section; specifically, the router must allow through any UDP traffic to ports greater than 1023 to the external DNS server. This means that even unsolicited requests could pass through. This could be addressed by reflexive access lists (as of Cisco IOS 11.3), but these require additional overhead and only one access control list may be applied to any one interface for ingress (and one for egress). Instead, the stateful firewall and DNS server are configured to add layers of security to address this hole.

Next, only selected ICMP traffic is allowed into the GIAC network. The permitted ICMP messages allow hosts on the internal network to receive replies to pings and error messages related to regular traffic such as destination unreachable and time exceeded messages.

Next, the honeypot traffic is allowed to pass through. The virtual IP hosting the HTTP and SMTP honeypot servers are filtered so that only the required ports (80,25) are allowed to pass through to them. All TCP traffic is allowed through to the tarpitted virtual IP addresses.

Finally, any traffic that does not match any of the explicitly permitted access control rules is blocked. In the event that a software bug disables this explicit deny all rule, then Cisco IOS implicitly denies all traffic (assuming at least one ACL has been applied).

Table 13: WAN Interface ACL

ACL	Comment
! Block spoofed and noise traffic	
access-list 101 deny ip 10.0.0.0 0.255.255.255 any	Private network
access-list 101 deny ip 172.16.0.0 0.15.255.255 any	
access-list 101 deny ip 192.168.0.0 0.0.255.255 any	
access-list 101 deny ip 224.0.0.0 31.255.255.255 any log	Multicast
access-list 101 deny ip 127.0.0.0 0.255.255.255 any log	Loopback
access-list 101 deny ip 169.254.0.0 0.0.255.255 any log	Local link block
access-list 101 deny ip 223.47.17.0 0.0.0.31 any log	Internal Network
access-list 101 deny ip host 0.0.0.0 any log	Traffic with no source IP address
access-list 101 deny ip 0.0.0.0 0.255.255.255 any log	Unallocated block
...	All unallocated blocks
! Permit VPN Traffic	
access-list 101 permit esp any host 223.47.17.30 log	VPN Traffic
access-list 101 permit udp any host 223.47.17.30 eq 500 log	IKE
! Permit DMZ Service Traffic	
access-list 101 permit udp any host 223.47.17.5 eq 53	DNS

access-list 101 permit tcp any host 223.47.17.6 eq 25	SMTP
access-list 101 permit tcp any host 223.47.17.7 eq 80	HTTP
access-list 101 permit tcp any host 223.47.17.7 eq 443	HTTPS/SSL
! Permit responses to outbound tcp traffic. Established keyword used here, but needs to also be addressed by stateful firewall	
access-list 101 permit tcp any host 223.47.17.6 gt 1023 est	Responses to outbound smtp
access-list 101 permit tcp any host 223.47.17.8 gt 1023 est	Responses to outbound http(s)
! Permit responses to dns queries ! This opens up possible unsolicited udp packets	
access-list 101 permit udp any host 223.47.17.5 gt 1023 log	DNS responses
! Permit only echo replies, all time exceeded and destination unreachable icmp traffic	
access-list 101 permit icmp any 223.47.17.0 0.0.0.31 0	ICMP Echo Reply
access-list 101 permit icmp any 223.47.17.0 0.0.0.31 11	ICMP Time Exceeded
access-list 101 permit icmp any 223.47.17.0 0.0.0.31 3	ICMP Destination Unreachable
! Permit honeypot traffic	
access-list 101 permit tcp any host 223.47.17.4 eq 80 log	Honeypot HTTP server
access-list 101 permit tcp any host 223.47.17.4 eq 25 log	Honeypot SMTP server
access-list 101 permit tcp any host 223.47.17.9 log	Tarpit traffic destined for 223.47.17.9 - 27.
access-list 101 permit tcp any host 223.47.17.10 log	
...	
! Explicitly deny all other traffic	
access-list 101 deny ip any any log	Deny and log all other traffic

VPN Interface Ingress

This interface only receives packets coming from the VPN gateway (with a source IP address of 223.47.17.30). The only two protocols allowed through this interface are IP ESP and UDP IKE. Logs for denied traffic include the MAC address in case something besides the VPN gateway is plugged into the VPN interface. The following table shows this access control list.

Table 14: VPN Interface ACL

ACL	Comment
access-list 102 permit esp host 223.47.17.30 any log	ESP traffic
access-list 102 permit tcp host 223.47.17.30 any eq 500 log	IKE
access-list 102 permit icmp 223.47.17.0 0.0.0.31 8 log	ICMP echo requests
access-list 102 deny ip any any log-input	Deny and log all other traffic

Firewall Interface Ingress

This interface receives packets from the DMZ or the firewall bound for the internet. The only TCP/UDP services allowed through this interface are HTTP, HTTPS, DNS and SMTP. ICMP echo requests are allowed out to satisfy the needs of the system administrators. No other ICMP traffic is allowed out

through this interface though.

Table 15: Firewall Interface ACL

ACL	Comment
access-list 103 permit udp 223.47.17.5 any gt 1023 log	DNS outbound requests
access-list 103 permit udp 223.47.17.5 any eq 53 log	DNS service
access-list 103 permit tcp 223.47.17.6 any gt 1023 log	SMTP outbound requests
access-list 103 permit tcp 223.47.17.6 any eq 25 log	SMTP service
access-list 103 permit tcp 223.47.17.7 any eq 80 est	HTTP Reverse Proxy
access-list 103 permit tcp 223.47.17.8 any gt 1023	HTTP(S) outbound requests
access-list 103 permit tcp 223.47.17.9 any eq 80 est log	Honeyd HTTP server
access-list 103 permit tcp 223.47.17.9 any eq 25 est log	Honeyd SMTP server
access-list 103 permit icmp 223.47.17.0 0.0.0.31 8	ICMP echo requests
access-list 103 deny ip any any log-input	Deny and log all other traffic

The VPN and firewall interfaces serve as the egress filters for the GIAC network. An additional WAN interface egress access control list could be implemented as an added layer of security, but this would add processing overhead for a redundant check. Instead, the firewall is relied upon for the extra layer of security.

Primary Firewall Security Policy

The primary firewall is PF running on an OpenBSD 3.6 platform. The firewall includes many rules that are redundant considering the filtering rules from the border router. The redundancy serves as an added layer of defense. The firewall also does packet normalization through PF's "scrub" rule. The packet normalization contributes to defense in depth by thwarting attacker's attempts at fingerprinting GIAC operating systems from outside the network. For the reader's information, PF processes rules in a last match wins fashion (exactly opposite of the Cisco access control list above). Additionally, an implicit pass all rule is added at the beginning of the rule base by PF (OpenBSD). Also by default, all traffic using IP options is blocked.

The 'quick' keyword is used in the PF filters to force a matched rule to be the last rule that PF compares the packet to. The 'keep state' keywords are used to enforce stateful inspection on the packets crossing the interface. The policy for stateful inspection is set by the "set state-policy if-bound" command in the PF configuration file. This means that when a state entry is created, the state is bound to the interface that it was created on. Packets that match a state in the state table must also be crossing the interface the state is bound to. The default behavior is not to bind the state to any particular interface. Another interesting thing about the stateful inspection is keeping state on UDP and ICMP messages. PF keeps state on these types of packets using a timeout

mechanism set by the “set timeout” command in the configuration file. If no matching packet is seen within the configured period, the state is dropped. Additionally, any ICMP messages with error messages related to any TCP or UDP traffic that has already created a state entry will pass in on that state entry without any specific rules for the ICMP traffic.

The first rule changes the default policy to deny packets unless explicitly permitted. The block policy for all rules is set to silently discard packets using the “set block-policy drop” command in the configuration file. Most all of the rules used in the PF configuration use macros and tables that would also be defined in the PF configuration file. Rather than enumerate all the macro and table definitions, intuitive names were chosen to aid in understanding the rules completely.

VPN Traffic

Next, the VPN traffic is locked down to only allow authentication, authorization and required application traffic is allowed through to the internal network. Each rule uses stateful packet inspection, including UDP traffic.

Border Router Traffic

Next, traffic coming from the border router is passed through. For traffic coming from the internet, only the traffic initiating transactions are enumerated on the external firewall interface. All other traffic should match a state defined by rules bound on the DMZ interface. Additionally, all traffic bound to the honeypot virtual web server, mail server or tarpits is allowed through.

DMZ Traffic

Traffic originating from the DMZ network may be going out to the internet or into the internal GIAC network. Stateful rules are used to create state table entries for packets coming in from other interfaces related to connections originally established from the DMZ network.

Internal Traffic

All traffic originating from the internal network heads either towards the DMZ network or to the VPN network except for SSH traffic from specified system administrator machines and ICMP echo requests. The SSH traffic allows the system administrators to remotely manage the routers, while the ICMP echo requests allow employees to see if hosts are alive on the Internet. In order to allow ICMP echo requests, the firewall must perform Network Address Translation (NAT) in order to allow internal hosts with private IP addresses to send routable ICMP echo requests. All traffic going through NAT must still pass the filter rules after translation in order to be let through by PF.

Table 16: Packet Filter Rule Base

Rule	Comment
------	---------

# Packet normalization	
scrub in all	Packet normalization
# Translation	
nat on \$ext_if from 192.168.1.0/24 to any -> (\$ext_if:0)	NAT addresses from the user network (still must be filtered)
# Change policy to default deny	
block log all	Change to default deny policy
# Traffic coming from vpn (source addresses will be the addresses from other side of bridge)	
pass in log quick on \$vpn_if proto tcp from any to <domain_controllers> port \$tcp_dc_ports keep state	Kerberos, LDAP and Netbios traffic required for domain controller functionality
pass in log quick on \$vpn_if proto udp from any to <domain_controllers> port \$udp_dc_ports keep state	Kerberos, LDAP and Netbios traffic required for domain controller functionality
pass in log quick on \$vpn_if proto tcp from any to <internal_mail_servers> port 25 keep state	SMTP traffic
pass in log quick on \$vpn_if proto tcp from any to any port 22 keep state	SSH traffic for remote system administrators
pass in log quick on \$vpn_if proto udp from any to <internal_dns_servers> port 53 keep state	Internal DNS traffic
pass in log quick on \$vpn_if proto udp from <vpn_gateway> to <syslog_servers> port 514	Syslog traffic
...	Add rules for all other GIAC application servers that remote users are required to access
# Traffic coming from the border router	
pass in log quick on \$ext_if proto tcp from <valid_external_ips> port > 1023 to <dmz_web_servers> port 80 keep state	HTTP requests
pass in log quick on \$ext_if proto udp from <valid_external_ips> port > 1023 to <dmz_dns_servers> port 53 keep state	DNS requests
pass in log quick on \$ext_if proto tcp from <valid_external_ips> port > 1023 to <dmz_mail_servers> port 25 keep state	Mail requests
pass in log quick on \$ext_if proto tcp from any to <honeypot> port {25, 80} keep state	Honeypot mail and web server
pass in log quick on \$ext_if proto tcp from any to <tarpits> keep state	Tarpit traffic
pass in log quick on \$ext_if proto udp from <border_router> to <syslog_servers> port 514	Syslog traffic
# Traffic coming from the DMZ	
pass in log quick on \$dmz_if proto udp from <dmz_dns_server> port > 1023 to <valid_external_ips> port 53 keep state	External dns requests
pass in log quick on \$dmz_if proto tcp from <dmz_smtp_server> port > 1023 to {<valid_external_ips> <internal_smtp_server>} port 25 keep state	SMTP traffic

pass in log quick on \$dmz_if proto tcp from <dmz_reverse_proxy> port > 1023 to <internal_web_servers> port {80 443} keep state	Reverse proxy traffic
pass in log quick on \$dmz_if proto tcp from <dmz_web_proxy> port > 1023 to <valid_external_ips> port {80 443} keep state	Web proxy traffic
pass in log quick on \$dmz_if proto icmp from any to <valid_external_ips> icmp-type 8 keep state	Outbound echo requests
pass in log quick on \$dmz_if proto udp from \$dmz_if/24 to <syslog_servers> port 514	Syslog traffic
# Traffic coming from internal network	
pass in log quick on \$int_if proto tcp from <internal_system_admins> to any port 22 keep state	System administrator SSH traffic
pass in log quick on \$int_if proto tcp from \$int_if/16 port > 1023 to <dmz_web_proxy> port 8080 keep state	Employees accessing internet through web proxy
pass in log quick on \$int_if proto tcp from <internal_mail_servers> port > 1023 to <dmz_mail_server> port 25 keep state	SMTP traffic from main server to relay
pass in log quick on \$int_if proto udp from \$int_if/16 port > 1023 to <dmz_dns_server> port 53 keep state	DNS requests for external resolution
pass in log quick on \$int_if from \$int_if/16 to \$vpn_if keep state	Internal users connecting to the VPN
pass in log quick on \$int_if proto tcp from any to <honeypot> port {25 80} keep state	Honeypot traffic
pass in log quick on \$int_if proto tcp from any to <tarpits> keep state	Tarpit traffic
pass in log quick on \$int_if proto icmp from 192.168.1.0/24 to any icmp-type 8 keep state	Echo requests

© SANS Institute

References

- Chen, Gao, Kwiat. "Modeling the Spread of Active Worms", 2003.
- Davie, Bruce S., Peterson, Larry L. Computer Networks: A Systems Approach. 2nd ed. Academic Press, 2000.
- "Entrapment." Legal Definitions Dot Com. n.pag. On-line. Internet. Feb. 2005. Available <http://www.legal-definitions.com/entrapment.htm>
- "Ethics." The Merriam-Webster Dictionary. 1994 ed.
- Haggard, Parks, Haggard & Lewis, P.A. "Notable Cases." n.pag. On-line. Internet. Feb. 2005. Available http://www.haggardparks.com/notable_cases.asp
- The HoneyNet Project. "HoneyNet Definitions, Requirements, and Standards." 1.6.0 (14 Oct. 2004): n.pag. On-line. Internet. Feb. 2005. Available <http://www.honeynet.org/alliance/requirements.html>
- The HoneyNet Project. "Know Your Enemy: GenII HoneyNets." (3 Nov. 2003): n.pag. On-line. Internet. Feb. 2005. Available <http://www.honeynet.org/papers/gen2/index.html>
- Internet Addressing Number Authority (IANA): "INTERNET PROTOCOL V4 ADDRESS SPACE" (Aug 3, 2004): n.pag. On-line. Internet. Mar. 2005. Available <http://www.iana.org/assignments/ipv4-address-space>
- Liston, Tom. "Tom Liston Talks about Labrea." n.pag. On-line. Internet. Feb. 2005. Available <http://labrea.sourceforge.net/Intro-History.html>
- OpenBSD. "PF User's Guide." n.pag. On-line. Internet. Mar. 2005. Available <http://www.openbsd.org/faq/index.html>
- Page Museum – LaBrea Tar Pits. "Frequently Asked Questions." n.pag. On-line. Internet. Feb. 2005. Available <http://www.tarpits.org/info/faq/faqfossil.html>
- Postel, J. "Transmission Control Protocol." RFC 793 (Sep. 1981): 85 pp. On-line. Internet. February 2005. Available <ftp://ftp.rfc-editor.org/in-notes/rfc793.txt>
- Spitzner, Lance. "Honeypots: Definitions and Values." (29 May 2003): n.pag. On-line. Internet. February 2005. Available

<http://www.tracking-hackers.com/papers/honeypots.html>

Tzu, Sun. The Art of War.

United States. Department of Justice. "Interception and Disclosure of Wire, Oral, or Electronic Communications Prohibited." 18 U.S.C. 2511 (20 April 2000): n.pag. On-line. Internet. Feb. 2005. Available HTTP: <http://www.cybercrime.gov/usc2511.htm>

© SANS Institute 2000 - 2005, Author retains full rights.