



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.

GIAC Certified Firewall Analyst (GCFW)
Practical Assignment
Version 4.0

John Creevey
28 August 2004

CONTENTS

1.	ASSIGNMENT 1 – FUTURE STATE OF SECURITY TECHNOLOGY	3
1.1	INTRODUCTION	4
1.2	VLAN AND IP BASICS	4
1.3	PRIVATE VLANs	6
1.3.1	Promiscuous ports	7
1.3.2	Community ports	7
1.3.3	Isolated ports	8
1.3.4	PVLAN Edge ports	8
1.4	VLAN ACLs	9
1.5	PRIVATE VLANs AND VLAN ACLs COMBINED	10
1.6	DEFENSE IN DEPTH	11
1.7	PVLAN AND VACL IMPACT	12
2.	ASSIGNMENT 2 – SECURITY ARCHITECTURE	14
2.1	INTRODUCTION	14
2.2	DESIGN REQUIREMENTS	15
2.3	NETWORK DESIGN	17
2.4	IP ADDRESSING	18
2.5	INTERNET ROUTER	19
2.6	FIREWALL DEVICES	20
2.7	VPN DEVICES	21
2.8	LOAD BALANCING DEVICES	23
2.9	IDS DEVICES	24
2.10	CISCO SWITCHES	26
3.	ASSIGNMENT 3 – FIREWALL POLICY	27
3.1	RULEBASE	27
3.2	FIREWALL OBJECT DEFINITIONS	28
3.3	RULEBASE EXPLANATION	29
3.4	RULEBASE ORDER	32
4.	REFERENCES	33

1. ABSTRACT

This assignment has one main aim. The main objective is to show that the author understands and has competency over the material taught throughout the GCFW course.

In order to do this, a number of assignment tasks have been issued. The first task asks the author to demonstrate some research and findings within a not so common area of Information Security and then relate these findings back to be relevant against the material taught during the course.

The second task provides a case study and set of requirements. Based on the information given the author must propose a network and security design that will fulfil the requirements. Again, emphasis must be placed on demonstrating that concepts learned throughout the course can be applied to a practical situation.

The last task involves outlining the basis firewall policy that could be applied to one of the firewalls used in the design.

Overall, I believe this assignment does infact demonstrate the material taught in the GCFW course.

© SANS Institute 2004, Author retains full rights.

Assignment 1 – Future State of Security Technology

1.1 Introduction

As outlined verbatim from the assignment specification, “The purpose of this assignment is to introduce or explore a particular technology or future technology that is not in common use in the Information Security Industry today, but that you feel will have a significant impact on perimeter security, or defense in depth”

The topic I have chosen for this assignment is “VLAN ACLs (VACLs) and Private VLANs (PVLANS) and their use in network defense”. I have chosen this topic from the list of available topics shown at http://www.giac.org/GCFW_wishlist.php.

In order to comply with the requirements of this assignment the following aspects of both PVLANS and VACLs will be explored:

What is the problem that these technologies mitigate?

How do these technologies work?

How does the usage of these technologies segment and restrict network traffic?

How do these technologies aid in network defense and defense in depth?

How can the technologies be used as a network defense tool?

The affect these technologies will have on the Information Security industry.

The affect these technologies will have on personnel tasked with the day-to-day operations of perimeter technology.

1.2 VLAN and IP Basics

In order to clearly show how VACLs and PVLANS can be used to aid in a defense in depth perimeter strategy, the reader must first adequately understand the concept of VLANs, how they work and their current use and constraints within typical DMZ environments.

The concept of a VLAN can be mapped to both layer-2 and layer-3 of the well-known OSI communications model. All hosts that share the same VLAN at layer 2 will also be members of the same IP subnet at layer-3. A VLAN is a set of hosts that all share the same broadcast domain.

For hosts in the same VLAN that wish to communicate with each other, they can communicate directly at layer-2. They use the ARP protocol to map a layer-3 IP address to a layer-2 MAC address and can speak to each other directly using only each other's MAC address. For hosts in different VLANs that wish to communicate, they can only talk to each other by transiting through one or more devices capable of routing their layer-3 traffic. Inter-VLAN traffic must be routed through a device or series of devices that have a common layer-3 routed protocol (IP, IPX etc.) configured on at least two of the router's interfaces.

Simply speaking, devices in the same VLAN can speak directly to each other without the need to go via any other layer 3 device. Devices in different VLANs however can only communicate via a layer-3 device such as a router or a firewall.

The problem with the operation of VLANs as described above is that typically, for a host communicating to another host within the same VLAN, there are no network access-control mechanisms that can be applied to this traffic. "Network access-control" is used to mean access-control that can be applied to traffic by a network-oriented device (router, firewall, switch etc.). It is of course always possible for a host itself to apply access-control through the use of desktop firewalls or other host based technologies.

Because VLANs operate in this manner, the only place that layer-3 access-control can be performed is on a layer-3 capable device that is routing inter-VLAN traffic. Traditional VLAN configurations offer no layer-3 access-control for intra-VLAN traffic.

The most obvious manifestation of this problem relates to the ability of a network to contain an intrusion or attack. If a hacker or malware compromises a host in a certain VLAN, normal VLAN operation allows them unrestricted network access to all the other hosts in the same VLAN. Under these circumstances, it's likely that the hacker or malware will also be able to attack and compromise the remaining hosts in that VLAN.

From the point of view of restricting the damage radius if a single host gets compromised, the normal operation of VLANs offers no protection to the devices in the same VLAN.

Another potential problem unable to be addressed by normal VLAN operation is the ability to restrict the type of network traffic sent between devices in the same VLAN. Suppose there is actually a valid requirement for devices in the same VLAN to communicate with each other. For instance, if a DMZ VLAN contains both a front-end webserver and a back-end application server, these two devices are required to be able to communicate with each other. Normal VLANs offer no way to restrict the type of traffic the hosts can send or receive because the communications do not go via a layer-3 device. Not being able to restrict this traffic can also aid an attacker in finding subsequent devices to compromise if an attack on a single machine is successful.

By solely relying on layer-3 devices in order to apply your network security policies, you are allowing attackers or malware extra room to move within your networks.

The use of VLANs also implies a degree of trust between all the devices on that VLAN. There may be circumstances where hosts must be placed on the same VLAN whereby this trust does not exist. This might occur for reasons of poor network or security design, lack of network scalability/flexibility or in the interests of IP address space conservation. Specifically, a service provider network that must house different customer devices on the same VLAN would fall into this category. Despite the reason why this situation may occur, traditional VLANs offer no solution to deal with the isolation of devices in the same VLAN that do not trust each other.

As will be described in the following sub-sections, the use of PVLANS or VACLs or a combination of the two can be used to combat all of the problems mentioned above.

1.3 Private VLANs

PVLANS are used to control the traffic that can be sent to and from hosts in the same VLAN. PVLANS are only effective at isolating hosts on a VLAN entirely away from each other. Depending on the exact PVLAN configuration, two hosts on the same VLAN either can or cannot communicate directly with each other.

PVLANS do not give you the ability to restrict the type of traffic that can propagate between hosts on the same VLAN. This function is provided by VACLs, which will be described in a subsequent section.

PVLANS operate on the concept of distinguishing between primary and secondary VLAN associations. When PVLANS aren't in use, a host will generally be a member of only one VLAN. The host VLAN membership defines the hosts' IP subnet, and also the other hosts that are on the same VLAN. In PVLAN terms, the primary VLAN is equivalent to the normal host VLAN in non-PVLAN terms. For PVLAN implementations, the secondary VLAN is associated with a host to define which other subset of hosts on its primary VLAN it is allowed to communicate with.

The IETF have drafted RFC 3069, which addresses the PVLAN concept. The IETF refer to the PVLAN concept as VLAN aggregation. In IETF terminology the primary VLAN is referred to as the super-VLAN. The secondary VLAN is referred to as the sub-VLAN.

The PVLAN concept that will be described within this assignment is based on the Cisco implementation. The Cisco implementation goes further and adds more functionality than what the IETF describe in RFC 3069. The IETF describe the use of sub-VLANs that are essentially used to create a VLAN within a VLAN. All hosts in a sub-VLAN can communicate with each other and their default gateway. Hosts on one sub-VLAN cannot communicate with hosts on a different sub-VLAN. The IETF make no attempt to designate ports as having different functions other than having membership with either super-VLANs or sub-VLANs.

Cisco, as well as using the primary and secondary VLAN concept, designates their switch ports to operate in a certain and different manner within the PVLANS themselves. The port designations used by Cisco are as follows:

- Promiscuous ports
- Community ports
- Isolated ports
- PVLAN edge ports

These port designations are described in the following sections.

The organisation security policy should define the access requirements of the hosts within your networks (DMZ or otherwise). If hosts or groups of hosts are not explicitly required to speak with other hosts, then the PVLAN concept should be used to restrict them. The traditional VLAN concept aims to logically group similar hosts together by placing them in the same VLAN. The PVLAN concept aims to further extend normal VLANs by grouping similar hosts within a VLAN to satisfy exact traffic flow restrictions and segregation requirements in an organisation's security policy.

1.3.1 Promiscuous ports

A host connected to a switch port that is configured in promiscuous mode is not restricted in any way as to the devices on its primary VLAN it can communicate with. It must be stressed that this only includes the primary VLAN. If hosts belong to different primary VLANs, then no matter what the PVLAN port designation, they will be unable to communicate unless they go through a layer-3 capable device. Thus, promiscuous switch ports will communicate with any other switch port so long as the primary VLAN is the same. This includes other promiscuous ports, isolated ports and community ports.

A switch port configured in promiscuous mode is suitable for connection to the default gateway of a VLAN. A host on a VLAN will (generally) always need to be able to communicate with its default gateway, and thus the promiscuous port is given access to communicate with any other port also in the same primary VLAN.

Alternatively, other devices that require connectivity to every host on a VLAN are also candidates for connection to a promiscuous port. Such devices might be network management stations or network backup hosts. As a rule though, it's better not to locate network management and/or backup services on the same VLAN as your hosts, especially if the hosts reside within an Internet facing DMZ environment.

1.3.2 Community ports

The community concept is essentially used to define a VLAN within a VLAN. The outer VLAN is the primary VLAN or in IETF terminology the super-VLAN. Each switch port no matter what the designation must belong to a primary VLAN. The inner VLAN, is referred to as the secondary VLAN, and defines the sub-set of hosts within the primary VLAN that are permitted to communicate with each other. Hosts within the same secondary VLAN are considered to be in the same "community" and thus their communications within the same secondary VLAN are not restricted.

Community ports are also permitted to communicate with promiscuous ports that belong to the same primary VLAN. It is essential that this is the case otherwise hosts in a secondary VLAN would not be able to communicate outside their subnet. Community ports are unable to communicate with community ports that belong to a different secondary VLAN and also isolated ports. This provides layer-2 isolation from other groups of hosts on the same VLAN.

Community ports are obviously useful when hosts in the same VLAN are required to talk to each other. An example of this might be for hosts that are clustered. Servers that are clustered for improved application performance often need to have interfaces in a common VLAN both from a layer-3 and layer-2 perspective. This is for the purpose of exchanging server state information so each server can be informed of the status of each other server in the cluster.

Alternatively, if you have devices that need to exist on a single VLAN for IP addressing purposes but also have different trust levels, then you can use community ports with different secondary VLANs to segment the hosts. Hosts of one trust level can be placed in a different secondary VLAN to separate them from hosts of a different trust level. A good example of this situation would be in a service provider environment. The service provider may house servers belonging to multiple customers on the same VLAN. If this is the

case, community ports configured for different secondary VLANs can be used to logically separate the servers of different clients.

1.3.3 Isolated ports

Hosts connected to isolated ports are only permitted to communicate with promiscuous ports that belong to the same primary VLAN. Normally this would only include the default gateway for a VLAN.

Isolated ports are not permitted to talk to a community port or other isolated ports. They provide the most restricted access of all the PVLAN port designations.

Isolated ports are valuable in a DMZ network environment where hosts do not generally need to communicate with other hosts in their same VLAN. Typically hosts in the DMZ will only communicate outside their IP subnet in order to provide services to the Internet, access services on the Internet, or act as a relay or proxy for services located in a different DMZ network or on the Internal LAN.

The use of isolated ports in this manner allows very granular host segmentation within a VLAN. As illustrated above, one of the main disadvantages of traditional VLAN use is that all hosts are accessible to all other hosts in that VLAN. This provides an easy avenue for attackers to jump around between different hosts in the same VLAN once they have compromised one of the hosts. The use of isolated ports will significantly restrict the visibility an attacker has in the VLAN once he has compromised a machine.

1.3.4 PVLAN Edge ports

PVLAN edge ports are similar in function to isolated ports. Cisco also refers to PVLAN edge ports as protected ports. Protected ports are relevant only to the switch on which they are configured. If a single switch contains two ports that are both configured as protected, traffic will not be permitted between the two protected ports. However, if protected ports are configured on different switches, then this provides no isolation at all between the ports.

In terms of ports on a single switch, traffic is permitted between protected ports and all other ports, however traffic is not permitted between two protected ports. Protected ports can only communicate with each other if they go via a layer 3 device.

Because PVLAN edge ports only have local relevance to the switch they are configured on, they differ from the other PVLAN port designations. So long as you trunk both primary and secondary VLANs between different switches, the normal PVLAN port designations can be used to secure a VLAN across a stack of switches.

The advantage of PVLAN edge ports or protected ports is that they are supported on the older and smaller range Cisco switches. This means they are suitable for use on legacy networks or smaller networks that are not built using the more high specification Cisco switches. This is specifically relevant for DMZ networks. Typically, DMZ networks might not be built using the larger module-oriented chassis-based switches such as the Catalyst 6000/6500 or 4000/4500 series switches. This may be done because the larger switches are cost prohibitive or simply because high port density is not required in a DMZ network. If this is the case, PVLAN edge ports might be the only type of PVLAN ports able to be configured on DMZ switches.

In actual fact, despite the higher cost, the modular Cisco switches make good candidates for use in a DMZ for a number of reasons. They do have higher port density, they offer the full spectrum of PVLAN functionality, they reduce the overhead of managing and monitoring multiple smaller switches and also they can integrate extra functionality in the form of firewall modules, intrusion detection modules and load-balancing modules.

1.4 VLAN ACLs

The primary purpose of VACLs is to filter the traffic that can be sent within a VLAN. Technically speaking, a VACL filters traffic as it enters the VLAN at layer-2. I.e., after it has been transmitted by a host and is entering the switch port.

VACLs can be used to apply access-control to traffic bridged within a VLAN or for traffic destined outside the VLAN in which it originated. Inter-VLAN traffic however is typically filtered on a layer-3 device as the traffic must transit this device to be routed to a different VLAN. Thus, it makes sense that inter-VLAN traffic is best filtered on the layer 3 device while intra-VLAN traffic is best filtered using a VACL.

This statement does make sense, however if you think about it, it might also be better to filter both inter-VLAN and intra-VLAN traffic using VACLs. Of course, intra-VLAN traffic can only be filtered by a VACL, so there is no option other than to use a VACL for this type of traffic. A layer-3 device however could potentially be used to filter inter-VLAN traffic once it has been bridged to the layer-3 device by the switch. The reason for not doing this is based on device performance. If you have unauthorised traffic flows that you would normally filter using ACLs on layer-3 devices, it makes better sense to filter them out using VACLs because you are removing the traffic from the network earlier. By filtering this traffic using a VACL rather than layer-3 ACL you are thus saving your switch the CPU cycles or processing overhead of having to bridge this traffic.

In actual fact, it makes total sense to configure both VACLs and layer-3 ACLs to filter this traffic. This provides advantages both of performance and security. The switch will no longer have to bridge traffic that does not comply with the security policy, and you also have implemented defense in depth by having the layer-3 device to fallback on if for some reason the VACL filtering does not work. The processing of VACLs is done in hardware at wire-speed with no performance decrease, so there is no performance drawback in using VACLs to process both inter-VLAN and intra-VLAN traffic.

Cisco VACLs can actually be used to filter more than just IP based traffic. VACLs can filter all of IP, IPX and MAC (layer-2) based traffic. This is simply done using normal IP, IPX or MAC based ACLs within the Cisco configuration.

The VACLs use normal Cisco ACLs to identify traffic. Once traffic has been matched against an ACL, an action can be taken to deal with the traffic. The actions that are available include drop, forward, log, redirect or capture.

Drop, forward and log are self-explanatory however redirect and capture allow the traffic to be sent to other ports on the switch so that packet capture devices can analyse, store or log the traffic. This might be required for archive and forensic purposes, or for logging and alerting purposes.

VACLs operate differently from conventional ACLs used on layer-3 devices. Specifically on Cisco routers, ACLs are applied to a specific interface in a specific direction. VACLs however are applied to the entire VLAN.

VACLs can be used to replicate the function of isolated ports in PVLANS. By configuring a VACL that only permits hosts in a VLAN to communicate with their default gateway this achieves the same end-result as an isolated port in a PVLAN. The difference is that isolated port configurations are only applied to that specific port, whereas VACLs are applied to the entire VLAN. This VACL would have the same effect as configuring every VLAN port as isolated, except of course the default gateway.

It's obvious to see then, that VACLs can be configured within a VLAN to restrict the type of traffic the hosts in the VLAN are permitted to send. This can also address the issues listed in the initial VLAN section above.

If you configure a VACL to prevent hosts in a VLAN communicating to other hosts in the same VLAN, then this significantly restricts the field of view an attacker has if he has compromised a host in that VLAN. This will prevent the attacker from being able to further compromise hosts in the same VLAN.

Also, a VACL can be used to stop a certain group of hosts in a VLAN from communicating with a different group of hosts. This addresses the issue where different trust levels exist between multiple hosts in the same VLAN. In order to do this however, you must be able to use normal ACL syntax to summarise the IP range used by the different groups of hosts in the VLAN.

It's clear then, that VACLs provide a very useful option when maintaining a defense in depth strategy and also give you extra options for implementing an organisation's security policy.

1.5 Private VLANs and VLAN ACLs Combined

PVLANS when used on their own provide good options for securing the traffic between hosts in the same VLAN. The same can be said of VACLs. When used in conjunction, they can complement each other to provide a very high degree of security for intra-VLAN traffic.

Some Cisco switches allow you to create and apply VACLs to both the primary and secondary VLANs used in PVLAN configurations.

Being able to create individual VACLs for secondary VLANs allows a very high level of access-control to be applied to intra-VLAN traffic. This however, is not possible across all the range of Cisco switches. On Cisco IOS based switches, VACLs are not applicable to secondary VLANs. If you apply a VACL to a primary VLAN, then this VACL also applies to all the secondary VLANs associated with that primary VLAN.

For Cisco switches that operate using certain versions of the Catalyst operating system (CatOS), it is indeed possible to configure and use different VACLs per secondary VLAN. For example, in CatOS versions 6.1(1) and above on the Catalyst 6500 series switch, it's possible to use different VACLs on different secondary VLANs.

When using PVLANS alone, there exists a chance that an attacker could send traffic between different isolated ports or different communities by routing it through the default gateway on the VLAN. In order for this to happen, the routing on the hosts must be altered to send intra-VLAN traffic to the default gateway. Alternatively, if the default gateway performs proxy-arp for the VLAN then this could subvert the security of PVLANS.

In order to ensure that PVLANS cannot be subverted in this manner, Cisco actually do recommend configuring VACLs to restrict traffic being sent where both the source and destination addresses belong to the IP subnet associated with the primary VLAN. So, in this instance, the VACLs are applied to the secondary VLANs however the IP subnet to filter actually represents the primary VLAN.

So, by using PVLANS to restrict the hosts that are able to communicate within the VLAN and using VACLs to restrict the actual type of traffic that can be sent between these hosts it gives incredible flexibility in your ability to implement your security policy correctly. In fact, by having these technologies available, it gives you extra options for when you are actually defining your security policy. Without these technologies, it's not possible to filter inter-VLAN traffic at all.

1.6 Defense in Depth

Defense in depth is the use of multiple technologies from multiple vendors in a layered/hierarchical approach to ensure that compromise of any single layer does not compromise the whole system/network.

PVLANS and VACLs aid greatly in creating a defense in depth security posture. Normally a DMZ will contain one or more filtering routers, one or more firewalls capable of stateful inspection and one or more application proxies for Internet services. This suite of technologies alone provides good security functionality from which to build a layered DMZ perimeter. By mixing and matching these technologies with the correct security policies and network design a formidable and secure network can be attained.

None of these technologies however address the VLAN issues that have been the discussed within this assignment. By adding PVLANS and VACLs to the list of available technologies that can be used to build your DMZ perimeter, you are increasing your ability to create and also maintain a defense in depth strategy.

It was demonstrated above that by using VACLs and router ACLs at the same time, this could prevent a failure or weakness of one technology opening up a hole in the network through which an attacker could compromise our systems.

If we add PVLANS and VACLs to our arsenal of security technologies, we have more options to choose from when defining our security policies. We also have more options to choose from when choosing how best to implement our existing security policies.

In this way, PVLANS and VACLs offer us extra security functionality and extra options to choose when formulating a network security design based on a defense in depth strategy.

1.7 PVLAN and VACL Impact

This section will aim to show the impact that PVLANs and VACLs might have on the Information Security industry. Also, the affect they will have on personnel tasked with maintaining them will also be explored.

I think that PVLANs and VACLs will have some impact on Information Security but not likely as large an impact as some other technologies like the firewall or IDS systems. As outlined in the defense in depth section, there already exists a large range of technologies available for securing our networks. The addition of PVLANs and VACLs simply gives us more choice.

In saying this though, PVLANs and VACLs do address a specific gap in the functionality provided by existing security products. There is no other technology that can provide layer-3 filtering of intra-VLAN traffic and for this reason I think the use and acceptance of PVLANs and VACLs will gradually become more and more prevalent across the industry.

This will be more so if Cisco continues their domination of the Ethernet switch market. Alongside this, if Cisco implements full PVLAN and VACL functionality in their lower range switches we will also likely see these technologies become much more commonplace in our DMZ networks.

I think at this point in time these technologies have not yet had a large impact on our industry. However, ultimately I think they will be introduced more and more and become quite standard in secure network deployments. I think eventually then, the impact of these technologies on how we design and implement secure network will be obvious.

In terms of looking at impact from a different perspective, you can also think about the impact these technologies will have on the personnel that need to maintain and manage them.

I think the main impact PVLANs and VACLs might have on personnel will be related to the areas of event logging, event alerting and also network support.

In terms of logging, VACLs now provide another source of logs that must be monitored for signs of malicious activity. If your Cisco devices already log to a central Syslog server then this simplifies the aggregation of the extra logging entries, however personnel still need to be able to interpret the logs and draw meaningful conclusions based on the logs that are generated by the VACLs.

The VACL logs may be logged centrally, correlated and alerted upon based on the detection of suspicious trends in the log entries. (Ie, similar to what an IDS would do) If this is the case, then the age-old problem of false-positives exists. You need to be able to tune the VACL logging and alerting so that only valid suspicious circumstances cause alerts to be generated. The impact of tuning the logging and alerting incorrectly could mean that personnel waste a lot of time investigating alerts coming from the VACL logs that do not turn out to be anything malicious.

Personnel responsible for supporting a network that uses traditional VLANs generally understand how they work and know how to troubleshoot these environments. For instance if there is a communications problem between two devices in the same VLAN, support personnel understand that there is nothing really between these devices other than a layer-2 switch.

By introducing PVLANS and VACLs you are changing the way that support personnel must understand the use of VLANs. Support personnel must have an intimate understanding of the configuration of the PVLAN and VACL security policies in order to understand how to troubleshoot issues with intra-VLAN communication. This might involve giving support personnel access to security documentation that details the security policies in order for them to understand the configuration. Ultimately, the impact here is by using PVLANS and VACLs you are changing the way that VLANs operate. Support personnel must understand this and have exactly knowledge of the policies otherwise they will be unable to troubleshoot and/or fix network issues in a timely manner.

Although the different impacts listed above do exist, I don't think they are prohibitive to the widespread adoption of PVLANS and VACLs within organisations. With adequate personnel training and preparation these issues and impacts will not stop organisations using these technologies. In the case for and against PVLANS and VACLs, the advantages certainly outweigh the disadvantages.

© SANS Institute 2004, Author retains full rights.

2. ASSIGNMENT 2 – SECURITY ARCHITECTURE

2.1 Introduction

Assignment 2 aims to take a given scenario and translate this into a functional and secure perimeter network design.

The scenario that has been given describes a company called GIAC Enterprises (GIACE). GIACE are a small company (50 employees) that market and sell fortune cookie sayings. GIACE use the Internet as their sole medium from which to sell their sayings. Most of the employees are located at the GIACE head-office, however a small number of employees are also located in four other sales offices distributed throughout different locations around the world.

GIACE are currently in the situation where they have a very small Internet presence. This small Internet setup is providing the basis for all of their online sales. Despite that GIACE only have a small Internet presence, they currently have an IT support group of 6-7 people. This IT support group aims to support both the Internet presence and also internal networks and systems. The management of GIACE have been thinking for some time that it's not cost effective to employ so many staff to support their systems. As a result, they have recently undertaken an internal review of their business processes, sales strategies and use of IT to support their business.

This review has yielded several significant concerns regarding how they are currently doing business. These concerns are listed below:

- The number of IT support personnel are too high and this is costing the business too much money to continue with this level of staffing.
- The current IT support personnel do not have the required networking and security expertise to adequately maintain the current Internet facing systems and infrastructure.
- Their current Internet presence and Internet facing systems will not scale to support the sales growth predicted over the next 6-12 months.
- GIACE management have heard through industry reports that hacker and worm/virus activity is currently prevalent throughout the Internet and still increasing. The management are concerned their Internet systems might not be secure and will be vulnerable to compromise. Because of the lack of skills of their current IT support, management do not have faith in the current IT staff to be able to address this issue.

In order for GIACE management to adequately address their concerns they have contracted the help of a reputable security services organisation. An organisation by the name "Security Inc" has been brought in to address the concerns listed above.

Security Inc have discussed the problems with GIACE management and proposed the following recommendations for change:

- The GIACE Internet presence and systems used to support their online sales needs to be redesigned and upgraded. The upgrade is required from the perspective of both systems performance and also security.
- GIACE should trim their IT support team to a group of 2-3 people. The people for the support team should be chosen based on their existing skill sets. The skill set of the entire team must be able to meet the support needs of the organisation.
- The remaining IT support team chosen should receive training in both networking and security technologies. This will enable them to better support their infrastructure.
- The management and maintenance of the newly designed Internet presence will be outsourced to a security managed service provider. This will enable GIACE to be able to support their systems with a reduced IT head-count. This will also enable the GIACE Internet facing infrastructure to be supported by personnel that have a high level of expertise and experience in networking and security.

2.2 Design Requirements

GIACE have developed the requirements for their Internet presence redesign. The requirements are listed as follows:

- The new Internet presence must support connectivity options for the following groups of users that collaborate with GIACE and use their systems:
 - Customers – Individuals and organisations that purchase fortune cookie sayings. Customers will access the GIACE front-end web infrastructure. The web-servers will provide the user a secure interface into the GIACE Internet purchasing system.
 - Suppliers – Organisations that provide fortune cookie sayings for GIACE to sell. Suppliers will access web-servers. These servers are specifically implemented for the purpose of facilitating B2B connectivity. Suppliers will connect to GIACE using a VPN service.
 - Partners – International organisations that translate and resell the cookie sayings. Partners will access web-servers. These servers will be the same B2B servers accessed by suppliers. Obviously the B2B interface will not be identical between suppliers and partners. The B2B interface and functionality provided will differ based on whether the user is a supplier or partner. Partners will also connect to GIACE using a VPN service.
 - GIACE Internal Employees – Users located at GIACE head-office. GIACE internal employees will connect to the existing LAN at head-office. They will have direct access to internal systems and also some of the systems in the DMZ. Internal users will be forced to use application proxy servers in the DMZ for access to the Internet.
 - GIACE External Employees – Remote-access users throughout the world. Employees not located at head-office will be based at one of four of the other GIACE workplaces around the world. Employees from these sites will

access head-office by using the VPN service. Employee access over the VPN will be limited to only those systems and services they need access to.

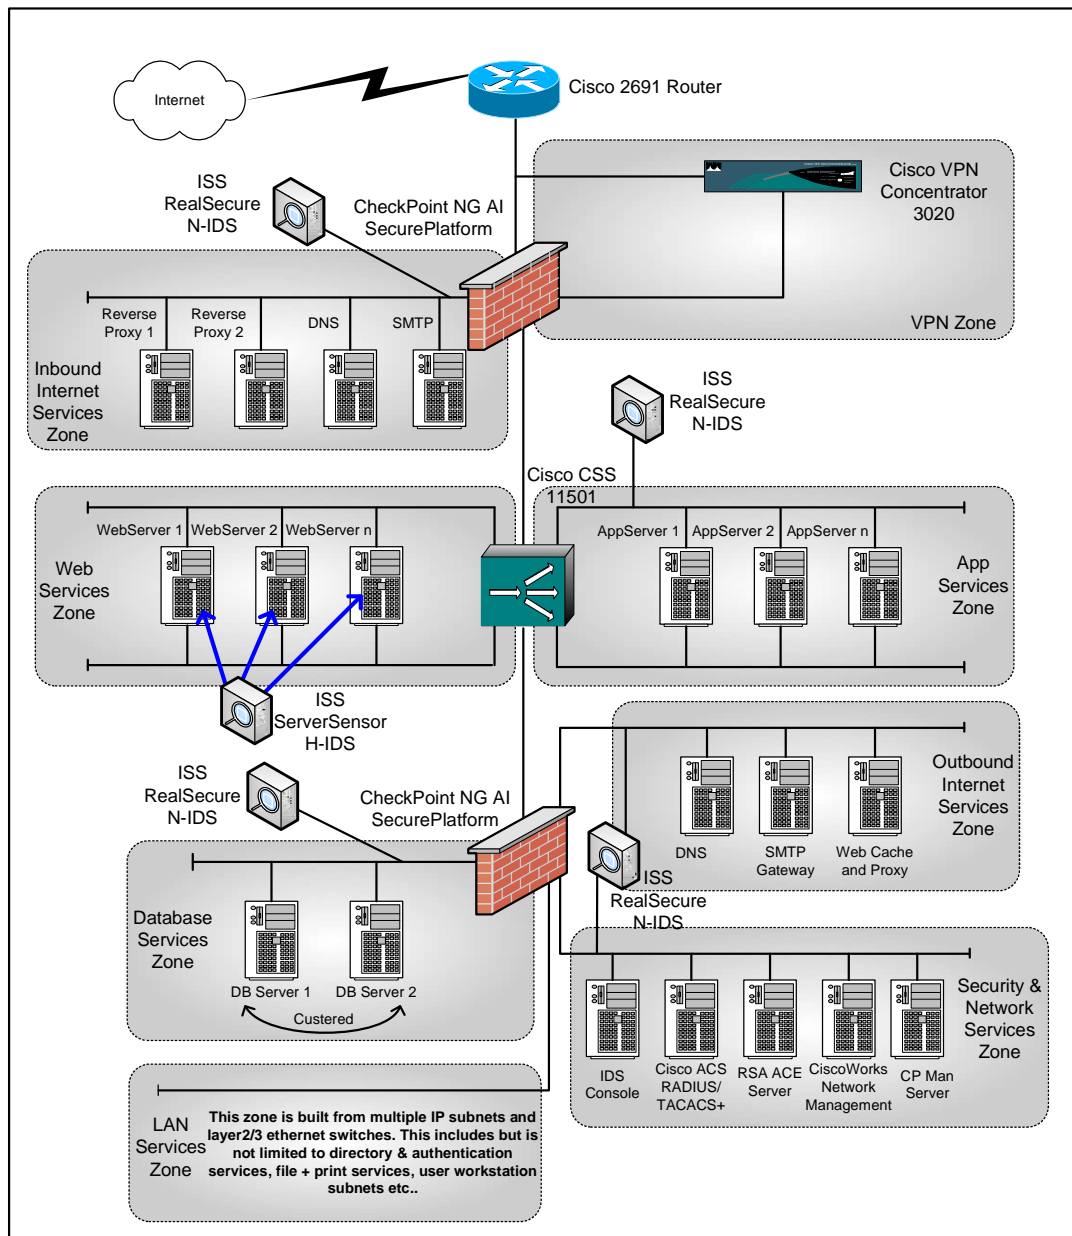
- The General Public – Individuals that access GIACE systems on the Internet for information. Those individuals not using the GIACE systems to purchase fortune cookie sayings can browse company information on the same web infrastructure that supports the purchasing system.
- Network and security design must be as simple as possible without sacrificing security.
- Network and security design should be as easy to maintain as possible. It's well understood that the hardest part of a security installation is with regards to its ongoing maintenance.
- The specific network components used in the design must:
 - Keep within budgetary constraints on capital expenditure.
 - Be cost effective. I.e Good performance and security ROI from the investment.
 - Be easy to maintain. I.e Not overly complex
 - Be able to be managed by the chosen managed services provider
- The specific network components to be used in the design are as follows:
 - Filtering routers, Ethernet switches, firewalls, VPN devices, IDS and load-balancing devices.
- There is no requirement for device level redundancy of any of the network or security infrastructure. Security Inc has done a resiliency costing analysis for GIACE. The results of this analysis indicate that the extra cost to duplicate devices and create a fully redundant architecture cannot be justified at this time. GIACE have asked Security Inc to focus on security and reliability in the design rather than resilience. GIACE have secured a 24x7x2 SLA and hardware support and maintenance contract with their managed services provider. This means that if any network component fails, the service provider will be onsite with a replacement in no more than 2 hrs. GIACE have weighed up this support cost against the cost of having a redundant architecture and it's shown to be more cost effective to have an aggressive support SLA rather than expend money upfront on extra capital purchases.

GIACE employ a relatively small number of people. In order to support the company's new initiative to rationalise their number of IT support staff, the decision has been made to outsource the management and maintenance of all the network and security components used in the new DMZ design.

Outsourcing in this manner reduces the burden on the newly downsized IT support staff. Another benefit of outsourcing this function is that it also ensures the maintenance and management of the new DMZ infrastructure is undertaken by personnel that has a high level of skills and experience in both network and security disciplines.

2.3 Network Design

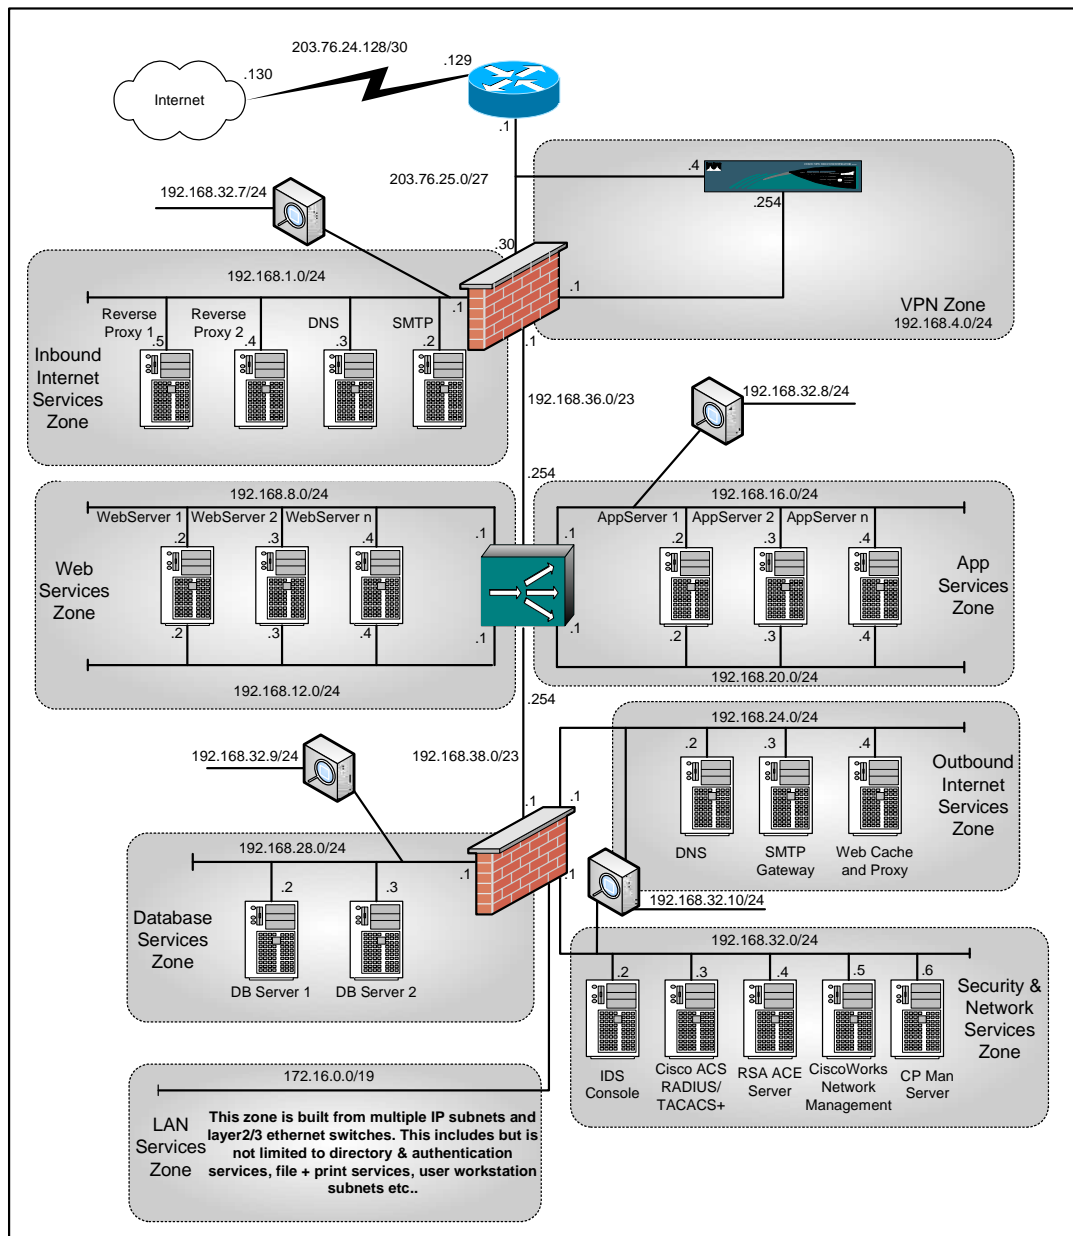
A simple DMZ network has been designed to address all the requirements set out by GIACE. The following diagram shows the topology of the proposed DMZ network design:



The design above has been developed to adhere to the design requirements. The design is quite simple yet also quite secure. The different parts of the DMZ perimeter have been broken up into functional zones. Each zone has been designed to fulfil a specific role in the network.

2.4 IP Addressing

The following diagram shows the IP addressing schema and individual addresses used as the basis of the GIACE perimeter redesign:



2.5 Internet Router

The Internet router chosen for this design is a Cisco 2691 series multiservice router. This router was chosen primarily on its ability to support the type of Internet connection in use by GIACE.

GIACE have a fractional E3 leased line connection to their ISP. GIACE find that this type of Internet connection best suits their business needs because they can easily upgrade the bandwidth.

The 2691 series router supports sub-rate T3/E3 network modules. If GIACE need to upgrade their Internet connection to full-rate T3/E3, then they would need to upgrade this router to at least a 3661, 3662 or 3725 series router. GIACE have confirmed that they cannot foresee the need for a full T3/E3 Internet connection and so the current router will be fine for at least a number of years.

The configuration and operation of the Internet router is important to the overall security architecture of the newly redesigned GIACE DMZ. The purpose of the router is more than to simply facilitate the GIACE Internet connection. Other functions to be provided by the Internet router are as follows:

- Ingress filtering on the WAN interface
- Ingress filtering on the Ethernet interface – some people prefer to do egress filtering on the WAN interface to achieve the same thing. Ingress filtering on this interface saves router resources by filtering the traffic before the router processes it.

The Internet router is used in this way to complement the security policy applied to the firewall. In terms of formulating a DMZ perimeter according to a defense in depth strategy, the Internet router is definitely the first line of defense.

The SANS GCFW courseware explains that the filtering ACLs on the router aims to block “absolutes” that may enter and leave the GIACE network. “Absolutes” is used to mean traffic that is not ambiguous. I.e., based on simply the source IP address, destination IP address or protocol you can be absolutely certain that you either do or do not want this traffic to enter or leave your network.

For example, the ingress ACL on the Ethernet interface would only allow traffic to pass that has a source address equal to a valid public IP address from the range allocated to GIACE. By blocking all other source IP addresses leaving your network, the GIACE network cannot be used to spoof IP addresses and you can attest to being a good “net neighbour”.

Likewise, ingress filtering is done on the WAN interface to block the same type of traffic. Traffic that you don’t want under any circumstances to enter your network. Traffic with RFC 1918 source addresses, or traffic known to be malicious from identifying the destination port can be blocked with confidence.

The filtering possibilities detailed above are not a complete list of the filtering you would undertake on your Internet routers. However, these are used as examples to show how you can filter traffic on the routers to complement the filtering and security policy on the firewalls. By doing this filtering on the routers you are reducing the amount and type of traffic the firewalls need to process.

Because of the actual function of this router, there are not too many options available for the placement of this device. The router must terminate the Internet connection and also connect to the same LAN as the external firewalls. The strengths of using a router in this capacity is that a static packet filter device is ideal for blocking this type of absolute traffic. The only potential weakness with the router doing ingress filtering in this manner is that the router has no stateful inspection capabilities. Thus, complex protocols cannot easily be securely filtered. By only filtering absolutes however, this is not a major issue.

2.6 Firewall Devices

Two layers of firewalls have been used in the new perimeter design for GIACE. Each layer of firewall adds specific security services to the overall security of the design. The two layers of firewalls are extremely important in the role they play within the perimeter defense in depth strategy.

The firewalls chosen for the design are based on CheckPoint FW-1 technologies. The exact version and platform of CheckPoint software will be CheckPoint NG AI R55 for SecurePlatform. Of course the latest hotfixes will also be applied.

CheckPoint has been chosen because it is a mature and well-established product in the market place. One of the advantages to using CheckPoint firewall software is the management products. CheckPoint provide a very comprehensive and intuitive suite of products for rule base creation and management, logging, software updates, license management and even performance monitoring.

SecurePlatform is an integrated package containing a securely optimised operating system and also the CheckPoint FW-1 software. The operating system is actually based on a cut-down and hardened linux distribution, thus making the platform both highly secure and also high performance. Because SecurePlatform is based on an open operating system, the cost of SecurePlatform is less than other platforms such as Nokia where you pay extra for the privilege of using the IPSO operating system.

SecurePlatform has been chosen for this design because it offers excellent performance and security for a very reasonable price tag.

The outer firewall is probably the most important element in the whole design. This firewall protects the entire DMZ from unauthorised access from the Internet. After the Internet router, this is the 2nd line of defense for the DMZ. The external firewall filters out most of the junk and only allows legitimate inbound connections to servers in the Inbound Internet Services Zone. The external firewall ensures that from the network perspective, connections into the DMZ are for valid traffic flows. It also ensures that traffic going further into the DMZ (ie into the web application) is also legitimate. If any of the servers in the Inbound Internet services zone are compromised, then the external firewall prevents the attacker from easily gaining access to other servers further on into the DMZ.

The web application used by GIACE has multiple tiers. The reverse proxy, front-end web-tier, 2nd level application tier and 3rd level database tier are all required to communicate to each other for the application to work correctly. By separating the application into tiers you can tightly control which application components can communicate to each other and how. By doing this you are isolating the most important part of the application (the data) furthest away from the attacker. This means the attacker would have to compromise reverse proxy, web server and application server in order to have un-restricted access to the

databases. This of course assumes there are no application level vulnerabilities that allow the attacker such access.

The different layers of firewalls aid in separating the application in this manner, and creating multiple layers of defense an attacker must penetrate before he can successfully compromise the entire application. In this way, each firewall adds to the defense in depth posture.

The inner firewall is also very important. The main function of this layer is to segment the internal LAN from the DMZ. Another very important aspect however is segmentation of the database tier away from the other front-end application tiers. This ensures that only valid requests are made to the databases from both inside the DMZ and also inside the LAN.

A very robust defense in depth strategy might use firewalls from different vendors as the internal and external firewalls. This is to combat against a vulnerability that might impact one layer of firewalls also impacting the 2nd or even third layer. In doing this though, it significantly increases the management overhead of maintaining the DMZ. Separate firewall management servers and access must be maintained, and certainly separate skill-sets must exist to be able to effectively manage such an installation. Security Inc made the decision that such a degree of defense in depth is not warranted for this design and the practicalities of ease of firewall management has meant firewalls on the same platform and from the same vendor have been chosen.

The strength of these firewalls is that they are a very established product and because they have evolved over a number of years the product is very manageable. As mentioned earlier, one of the main stumbling blocks to a perimeter network is not the design or installation itself but rather the ongoing maintenance and management. The best design in the world will be useless if it's not maintained according to sensible security policies and practices. The management tools available to manage CheckPoint firewalls are very good and thus, this makes these firewalls easier than others to maintain in the correct way. Thus, in the long run, your security policy will stay in-check and the risk of mis-configurations or policy violations is reduced.

Thus, by using two layers of firewalls in this manner you get very granular control over the traffic that is entering and exiting the DMZ perimeter. This has a two-fold effect. It keeps nosey people on the Internet and also on the internal LAN away from the critical application and database infrastructure in the DMZ. It also ensures that if people do poke around they have limited visibility as to what they can access.

2.7 VPN Devices

The VPN device chosen is a Cisco 3020 VPN Concentrator. This has been chosen on the basis of performance, functionality and manageability. The performance of this device is good as it does its crypto processing in hardware and can support 50 Mbps of encrypted throughput. The next smaller device than this one is the 3015 concentrator, however this device can only support 4Mbps of encrypted throughput. 50 Mbps is far more than GIACE will likely ever need, however with many concurrent sessions from suppliers, partners and remote GIACE offices, 4Mbps is not enough.

From a functionality perspective, the Cisco VPN concentrator provides all the VPN options that are required. The VPN security policy to be applied to GIACE VPN traffic is as follows:

- Partners – LAN-to-LAN - ESP-AES128-SHA-1, IKE-preshared
- Suppliers – LAN-to-LAN - ESP-AES192-SHA-1, IKE-preshared
- GIACE remote users – LAN-to-LAN - ESP-AES256-SHA-1, IKE-certificates

The different security profile that will be applied to each type of traffic has been based on the sensitivity of that traffic. GIACE remote users will be transmitting sensitive confidential information, and thus it was deemed this warrants a higher level of encryption than communications to partners and suppliers. Likewise, suppliers that upload confidential pricing details and product information are deemed to have a greater security need than partners that just download the latest fortunes that can be sold.

The concentrator supports both pre-shared secrets and certificates as a means to authenticate IKE phase-1 connections. Because relationships with partners and suppliers are more transient than the connections to GIACE remote users, this authentication method is chosen to use pre-shared keys. Because the GIACE remote user offices however are static, the remote-office infrastructure and head-end VPN concentrator mutually authenticate using X.509 based certificates. In this instance, because GIACE are a very small organisation, an external certificate authority will create and provide the certificates.

Thus, it can be seen that the Concentrator supports all the required IPSEC security services and protocols that contribute to creating and maintaining secure tunnels. Such other parameters such as IKE SA and IPSEC SA lifetimes will also be tuned depending on the sensitivity of the data traversing each tunnel.

One current limitation of the VPN concentrator is the lack of support for AH. AH in tunnel mode provides authentication over the entire packet, excluding the mutable fields. This provides for the highest level of packet authentication and integrity services. It is assumed that an implementation of AH on the concentrators will be forthcoming with subsequent software releases and AH will be used for GIACE remote user connections when this is available. For this reason, the public interface of the concentrator is not connected directly to the firewall. The public interface of the concentrator is assigned a publicly routed Internet IP address. This is to enable AH communications when the functionality is released. It's widely known that AH has problems through NAT devices and if the public IP address of the concentrator was being NAT'd by the external firewall then AH would not be possible.

Because the public interface of the concentrator is not protected by the firewall, filters on the concentrator will be configured to only accept traffic that is absolutely necessary for IPSEC tunnel creation and management.

The inside interface of the concentrator connects directly to the external firewall. Thus all traffic that enters the GIACE DMZ from one of the supplier/partner/GIACE VPN tunnels will be required to pass through the firewall. This will enable GIACE to restrict access using the firewall policy. Users at the other end of the VPN tunnels are requested to NAT their traffic behind nominated IP ranges. Thus, the firewall policy can differentiate between VPN users by their source IP address and apply an appropriate access policy for that user.

If required, the firewall could also apply an authentication requirement to the services that are accessed by each of the third parties. This would apply a second level of authentication over and above the authentication to actually establish the VPN. This authentication could be back-ended to the Cisco RADIUS/TACACS+ server or even perhaps the ACE Server if very strong authentication is required. This would adhere to the defense in depth principle by having multiple forms of authentication.

Also important aspect in the security of the VPN tunnels is the ability to route traffic directly between the tunnels without going to the firewall. Network lists are defined on the concentrator that detail the source and destination addresses of traffic that is permitted down each of the tunnels. Because the network lists for each tunnel will include a unique range of source IP addresses, there is no chance that inter-tunnel routing can occur. If one party did attempt to route traffic into another tunnel, the configured tunnel network lists would prevent this and the traffic would simply be dropped. This is an example of the VPN concentrator itself adding to the overall security of the VPN proposal by securing the traffic flows into and out of the VPN. The firewall cannot secure this traffic so the concentrator adds defense in depth to the overall security.

Thus, it is evident that the concentrator provides filtering, authentication, encryption and routing services alongside the filtering, authentication and NAT functionality provided by the firewall. Together as an efficient complement to each other they act to secure all the traffic entering and leaving the DMZ via the VPN service.

2.8 Load Balancing Devices

Central to the design and operation of the GIACE DMZ network is a Cisco Content Services Switch (CSS) model 11501. The Content Services Switch is another name for what is commonly known as a load-balancer.

The load-balancing device is literally deployed in the centre of the DMZ network. This is not done for security purposes, rather performance and service availability. Obviously, it's critically important that the DMZ infrastructure is secure and deterministic, however the actual reason for the DMZ in the first place is to facilitate access to web based applications and other services. By deploying the load-balancer in this manner, any number of servers in any part of the DMZ can be load-balanced. This provides a high level of application scalability and performance for DMZ applications.

Where it is required for service access, the CSS provides load-balancing services. Where load-balancing is not required, the load-balancer simply acts as a router from the inside of the DMZ to the outside or vice versa.

The GIACE web presence used to sell fortune cookie sayings relies on the majority of the infrastructure in the DMZ. Specifically the web-servers in the web services zone, and also the application-servers in the App services zone. When users want to buy fortune cookies, they connect to <http://www.fortunecookies.com>. This translates to a public IP address which is NAT'd by the external firewall to one of the reverse proxy servers in the Inbound Internet services zone. From here the reverse proxy servers connect to the front-end webserver using a pre-configured virtual IP address (VIP). This VIP represents the webserver farm in the Web services zone. The load balancer accepts inbound connections on the VIP address and distributes the connections to each of the webserver based on the configured load-balancing algorithm. The webserver then need to talk to

the application servers. The web servers connect to a single VIP, which is managed by the load-balancer and represents the application-server farm. The connection load is then distributed to the available application servers. In turn, the application servers access the database servers on a VIP presented to them and managed by the load-balancer.

It's obvious then that the load balancers provide essential services to the DMZ. In the primary instance the main application being load-balanced is the sales portal for the GIACE fortune cookie sayings. However, any service supported by the CSS could be load balanced offering a great deal of flexibility to how services are offered within the DMZ.

In terms of supporting a defense in depth strategy, the CSS boxes don't really offer a lot of defense capabilities. At the same time however, their use does not sacrifice any security functionality. By placing the load-balancer in between the internal and external firewalls you can use the firewall policy to control exactly which users get access to the configured VIPs, which in turn controls user access to DMZ services.

2.9 IDS Devices

A combination of both host and network IDS sensors have been used in the GIACE perimeter design. In most cases network sensors have been used because network sensors can be deployed more easily within a network segment to inspect all the traffic within that segment. Host sensors have been used where network sensors are not practical. For instance, in the Web services zone, you are not able to place a network sensor in front of the webserver farm. This is because the majority of traffic hitting the web servers will be http traffic encrypted within SSL sessions. Thus, a network sensor will not be able to inspect an encrypted http data stream. For this reason, host sensors will sit directly on each of the web servers and have the ability to inspect the sessions after they have been decrypted out of their SSL session.

Network IDS sensors have been placed at strategic places throughout the newly designed DMZ network. The placement of these devices has been done in such a way as to have the best chance at catching an attempted known exploit being carried out.

In his book, Stephen Northcutt discusses the relative merits of IDS placement around your network perimeter. Stephen talks about placement inside or outside your external firewall and why you might do one or both of these options. The GIACE IDS design does not include a network sensor outside the external firewall. This network sensor is obviously going to see every attack executed on the GIACE DMZ, which could potentially flag many alerts and cause a lot of work for the analyst monitoring the IDS console. Rather than take this approach and log "attacks", GIACE will place their IDS network sensors inside the external firewall and rather focus on logging and alerting "intrusions". Again, this concept has been taken from Stephen's book.

Most of the zones in the GIACE DMZ have IDS sensors placed inside. The reason for this placement strategy is so that the IDS policy on each sensor can be more accurately tuned to prevent a high occurrence of false positive alerts.

For example, the IDS sensor in the Web services zone would be tuned to look for exploits aimed at web-servers and ftp servers. This enables you to focus your view into the data streams and concentrate on searching for exploits that are actually relevant to the data stream being analysed. There is no point searching for SMTP exploits in a data stream that will never contain SMTP for instance.

This concept has been extended across all the zones in the DMZ. Obviously, the Inbound Internet services DMZ houses multiple services of different types, so this IDS policy will be much focused on being able to identify and alert on exploits for all the servers and services housed in this zone.

No sensor has been placed in the VPN zone, because the placement of sensors in all the other zones will cover all the traffic originating from this zone.

The actual product set used for both the host and network IDS sensors is based on the suite of IDS products provided by ISS. The versions proposed are RealSecure 10/100 version 7.0 for the network sensors and ServerSensor version 7.0 for the host sensors. Both of these products are the latest versions of the software offered by ISS and thus represent the latest developments in IDS technology from the perspective of both host and network IDS analysis.

Alongside the actual IDS sensors themselves, the ISS SiteProtector package will be installed along with the complementary SecurityFusion module. SiteProtector is the centralised management platform for the IDS sensors providing IDS event management, alerting, policy definition and deployment, automated signature file updates and an extensive reporting interface into the event database. SecurityFusion is an additional module added on top of SiteProtector that provides event correlation and prioritisation.

The reason that ISS has been chosen as the IDS product is because of its track record, maturity and evolution as a leading technology in the IDS industry. The RealSecure and ServerSensor products are extremely capable as IDS sensors in their own right. However, as equally important as this is very functional management interface.

It is well understood that the IDS offering from ISS is certainly a lot more expensive than other IDS options such as snort. However, the GIACE management stresses that the product sets they use must be as easy to maintain as possible. Thus, even though snort is much cheaper, due to its open source nature, it's also more complex to maintain and update. ISS and their management products make the actual ongoing maintenance and use of their IDS products easy to reduce the burden on operational staff. In the eyes of GIACE management, this benefit is worth paying for.

The network sensors will be installed into each network segment by using a network tap. This will ensure the sensors can analyse the entire full-duplex data stream for the segment. This also removes the need to span switch ports. Spanning switch ports on busy segments can increase switch CPU utilisation. By actually using passive network taps, this removes this potential problem.

The detection interface of each sensor will be deployed in stealth mode on each segment so as to protect the sensor itself from attack. The sensor management interfaces will all connect back to the Security & Network Services zone. Over this local network, they will send events and logging to the SiteProtector management console. Because the management console and sensors are effectively directly connected to the same LAN, this means that even if there are network connectivity problems in certain parts of the DMZ, IDS events and alerts will continue without interruption.

It can be seen from the placement of the IDS network devices and placement of the IDS host devices where the network devices will not be effective that this indeed prescribes to the defense in depth strategy. Firewalls and other such technology are not 100% perfect. Although you can design a secure DMZ that will greatly mitigate against the likelihood of a successful attack, it's impossible to say that compromise will never occur. Thus, the

various IDS devices are that last line of defense that will alert to the presence of malicious activity if it does somehow penetrate the various other lines of defense in place.

2.10 Cisco switches

Each zone will be physically facilitated by one or more Cisco 3550 series Ethernet switches. Cisco are industry leaders in Ethernet switching technology and these switches have been chosen because they are well-known and will be easy to support and maintain. Again, GIACE management have stressed this is an important facet of all of the equipment that is purchased.

Not only are these switches relatively easy to manage, they also provide a high level of performance for DMZ based switches. They do not have the large switching backplanes of their larger counterparts such as the 4000/4500 series however they have more than enough port capacity and backplane switching capacity to be functional within this design.

Generally, each set of switches in each zone will be configured with one layer 2 VLAN that maps that zone. The Web services and App services zones however are an exception to this. For these zones, the layer 3 VLANs are routed using interfaces on the CSS load-balancer. For each of the Web services and App services zones there will be two VLANs used. One VLAN is an ingress VLAN that incoming connections are load-balanced across and the other out an egress VLAN that outgoing connections are routed across.

This routing/VLAN design has been done for traffic control purposes. Rather than have a webserver or application server with an interface to send and receive data in one VLAN, each server has two interfaces. Connections are received on one interface and transmitted on the other. Routing is never enabled on the servers between the VLANs. This provides good traffic control and routing functionality through the load-balancer.

The webserver and application servers have no requirement to communicate amongst each other. Thus, in order to lock down the servers within these local VLANs, access control in the form of VACLs will be applied to these VLANs. The VACLs will have two purposes.

The first purpose will be to ensure that the servers do not communicate with each other. In order to achieve this, a MAC based extended ACL will be used within the VACL to ensure the servers can only communicate with the MAC address of their default gateway.

The second purpose is to ensure the traffic flows through the servers as expected. In order to achieve this, IP extended ACLs will also be used within the VACL to make sure that traffic cannot be initiated and routed out the ingress VLAN and that traffic cannot be routed in the egress VLAN.

The VACLs will help ensure that the traffic transmitted by the servers is in accordance with the security policies and also what is expected from normal operation of the applications.

Again, this aims to re-inforce the defense in depth principle. The VACLs provide access control to complement the stateful inspection done on the firewalls. The VACLs also ensure that if a local webserver or application server is compromised that the attacker cannot also compromise one of the adjacent machines also.

3. ASSIGNMENT 3 – FIREWALL POLICY

3.1 Rulebase

The firewall policy below is used by the external firewall. This firewall policy will provide connectivity as per the design requirements section in Assignment 2. The FW-1 implied rules have been disabled and explicit rules have been created to cater for this. This makes the firewall policy more secure.

Rule	Source	Destination	Service	Action	Track
1	CP-Man-Server Firewall-Ext	Firewall-Ext CP-Man-Server	Firewall1	Accept	Log
2	CiscoWorks	Firewall-Ext	snmp-read echo-request SSH	Accept	Log
3	Any	Firewall-Ext	Any	Drop	Log
4	Firewall-Ext	CiscoWorks	snmp-trap syslog	Accept	
5	Not(GIACE-Nets)	Public-SMTP	SMTP	Accept	Log
6	DMZ-Ext-SMTP DMZ-Int-SMTP	DMZ-Int-SMTP DMZ-Ext-SMTP	SMTP	Accept	
7	DMZ-Ext-SMTP	Not(GIACE-Nets)	SMTP	Accept	
8	Not(GIACE-Nets)	Public-DNS	domain-udp	Accept	Log
9	DMZ-Int-DNS	DMZ-Ext-DNS	domain-udp	Accept	
10	DMZ-Ext-DNS	Not(GIACE-Nets)	domain-udp	Accept	
11	Secondary-Internet-DNS	Public-DNS	domain-tcp	Accept	Log
12	Any	Public-RP1 Public-RP2	http https	Accept	Log
13	DMZ-Ext-RP1 DMZ-Ext-RP2	Web-Tier-VIP	http https	Accept	Log
14	Remote-Offices-Source-NAT	GIACE-Nets	Any	Accept	Log
15	CiscoWorks	VPN-Private	https snmp_read	Accept	

			echo-request		
16	Any	Any	Any	Drop	Log

3.2 Firewall Object Definitions

Object Name	Object Type	IP Address	Mask
Firewall-Ext	Firewall-1	192.168.36.1 203.76.25.30 192.168.4.1 192.168.1.1	255.255.255.255
CP-Man-Server	Management Station	192.168.32.6	255.255.255.255
CiscoWorks	Workstation	192.168.32.5	255.255.255.255
DNS-Int	Workstation	192.168.24.2	255.255.255.255
DNS-Ext	Workstation	192.168.1.3	255.255.255.255
GIACE-Nets	Network	172.16.0.0 192.168.0.0	255.255.224.0 255.255.192.0
Public-SMTP	Workstation	203.76.25.2	255.255.255.255
DMZ-Ext-SMTP	Workstation	192.168.1.2	255.255.255.255
DMZ-Int-SMTP	Workstation	192.168.24.3	255.255.255.255
Public-DNS	Workstation	203.76.25.3	255.255.255.255
DMZ-Ext-DNS	Workstation	192.168.1.3	255.255.255.255
DMZ-Int-DNS	Workstation	192.168.24.2	255.255.255.255
Secondary-Internet-DNS	Workstation	203.76.28.17	255.255.255.255
Public-RP1	Workstation	203.76.25.5	255.255.255.255
Public-RP2	Workstation	203.76.25.6	255.255.255.255
DMZ-Ext-RP1	Workstation	192.168.1.5	255.255.255.255
DMZ-Ext-RP2	Workstation	192.168.1.4	255.255.255.255
Web-Tier-VIP	Workstation	192.168.36.5	255.255.255.255
Suppliers-VPN-	Network	192.168.40.0	255.255.252.0

Source-NAT			
Partners-VPN-Source-NAT	Network	192.168.44.0	255.255.252.0
Remote-Offices-Source-NAT	Network	192.168.48.0	255.255.252.0
VPN-Private	Workstation	192.168.4.254	255.255.255.255

3.3 Rulebase Explanation

This section explains the firewall rules that have been used and why.

1	CP-Man-Server Firewall-Ext	Firewall-Ext CP-Man-Server	Firewall1	Accept	Log
---	-------------------------------	-------------------------------	-----------	--------	-----

This rule is to allow the CheckPoint management server to talk to the firewall and also the reverse for the required communications protocols. CheckPoint uses a group of proprietary protocols for firewall management functions like policy download, logging etc. These functions are contained with the service group called Firewall1. Rather than list the services individually in the rulebase, the required extra services have been added to this group for tidiness.

2	CiscoWorks	Firewall-Ext	snmp-read echo-request SSH	Accept	Log
---	------------	--------------	----------------------------------	--------	-----

This is to enable the CiscoWorks management console to access the firewall on management protocols. Although Ciscoworks is generally a platform only used to manage Cisco devices, it can be used very generally to undertake icmp polling and snmp-trap reception for non-Cisco devices. Because of the large amount of Cisco devices in the design, Cisco was proposed for managing them and the non-Cisco devices can still be partially managed. These protocols are required for that management.

3	Any	Firewall-Ext	Any	Drop	Log
---	-----	--------------	-----	------	-----

This is the standard firewall stealth rule. Drop anything not required for firewall management. Prevents unauthorised access to the firewall itself. It is important to log this access.

4	Firewall-Ext	CiscoWorks	snmp-trap syslog	Accept	
---	--------------	------------	---------------------	--------	--

This permits the firewall to talk back to the Ciscoworks server. SNMP traps will be sent back along with syslog. This rule needs to be explicitly allowed because the explicit drop rule at the end of the rulebase block all other traffic originating from the firewall. This prevents the firewall being compromised and used as a launch pad for other attacks.

5	Not(GIACE-Nets)	Public-SMTP	SMTP	Accept	Log
---	-----------------	-------------	------	--------	-----

This allows users on the Internet to connect to the external SMTP mail gateway. The Public-SMTP is the public Internet address that will be statically NAT'd by the firewall.

6	DMZ-Ext-SMTP DMZ-Int-SMTP	DMZ-Int-SMTP DMZ-Ext-SMTP	SMTP	Accept	
---	------------------------------	------------------------------	------	--------	--

This allows the SMTP gateway in the external DMZ to talk to the SMTP gateway on the internal DMZ. This allows for excellent separation of SMTP connectivity between the Internet and the e-mail gateway on the LAN. At each stage the format and integrity of the SMTP message and payload can be verified. Also allows SMTP gateway on internal DMZ to communicate with SMTP gateway in external DMZ, to deliver mail to the Internet.

7	DMZ-Ext-SMTP	Not(GIACE-Nets)	SMTP	Accept	
---	--------------	-----------------	------	--------	--

This allows the SMTP gateway in the external DMZ to deliver mail to the Internet. Because rule number 6 is above this rule, it stops delivery of mail to any internal systems other than the other SMTP gateway in the DMZ. This is a case where rulebase order is important.

8	Not(GIACE-Nets)	Public-DNS	domain-udp	Accept	Log
---	-----------------	------------	------------	--------	-----

This allows Internet DNS servers to access the public Internet address of our DNS server. Our DNS server is authoritative for the GIACE domain, and thus must be able to return authoritative DNS records for this zone. The Public DNS IP address will be NAT'd by the firewall to the real address of the DNS server within the Inbound Internet services zone.

9	DMZ-Int-DNS	DMZ-Ext-DNS	domain-udp	Accept	
---	-------------	-------------	------------	--------	--

Allows the internal DMZ DNS server to resolve Internet DNS names using the external DMZ DNS server. Can return authoritative records for GIACE domain or do recursive queries for other Internet domains.

10	DMZ-Ext-DNS	Not(GIACE-Nets)	domain-udp	Accept	
----	-------------	-----------------	------------	--------	--

Allows the external DMZ DNS server to access other DNS servers on the Internet for resolution.

11	Secondary-Internet-DNS	Public-DNS	domain-tcp	Accept	Log
----	------------------------	------------	------------	--------	-----

Allows our ISP, who maintains secondary DNS records for the GIACE domain (fortunecookies.com) to do zone transfers of the authoritative zone when the records have changed.

12	Any	Public-RP1 Public-RP2	http https	Accept	Log
----	-----	--------------------------	---------------	--------	-----

Allows anyone (including internal users and also Partners and Suppliers that come through the VPN) to access the reverse proxies for the purpose of accessing the web application.

13	DMZ-Ext-RP1 DMZ-Ext-RP2	Web-Tier-VIP	http https	Accept	Log
----	----------------------------	--------------	---------------	--------	-----

This permits the reverse proxies to have access to the web servers in the Web services zone. Notice that only access is given to the Web-Tier-VIP IP address. This is then load-balanced and connections are distributed between the available web servers.

14	Remote-Offices-Source-NAT	GIACE-Nets	Any	Accept	
----	---------------------------	------------	-----	--------	--

Allows the employees located in the remote offices access back to the corporate LAN. Because the VPN and firewall infrastructure is new, the GIACE management have requested that the remote employees have full and open access to the entire network. Security Inc has advised against this full and open access for the obvious security implications. Security Inc have agreed with management to turn logging on this rule and analyse the applications and servers accessed by the remote employees for the purposes of tuning the firewall policy and locking down the services and servers to be as granular and as secure as possible. GIACE management have agreed with this.

This is the only reason this rule has been left in this state. Allowing "Any" access is a bad idea for the long-term.

15	CiscoWorks	VPN-Private	https snmp_read echo-request	Accept	Log
----	------------	-------------	------------------------------------	--------	-----

Allows the CiscoWorks management platform to communicate with the VPN concentrator for management. The concentrator is a Cisco device and can integrate easily into

CiscoWorks. It's managed and configured by an embedded web-server, hence the https access is required to the private interface.

16	Any	Any	Any	Drop	Log
----	-----	-----	-----	------	-----

This is the explicit drop. Normally firewalls operate based on the convention of "that which is not expressly permitted is denied". This last explicit drop rule is the rule that denies everything that is not permitted above.

3.4 Rulebase Order

The CheckPoint firewall compiles a list of rules from top to bottom. When the rulebase is being processed, the rules are sequentially analysed and matched against the received traffic. Because the rules are analysed sequentially and top to bottom, the order in which you place your rules can be very important.

Take for instance rule number 7. This rule permits the SMTP gateway in the external DMZ to deliver mail to the Internet. Rule number 6 is also important. This permits the SMTP gateway in the external DMZ to deliver mail to the SMTP gateway in the internal DMZ. If I had not located rule number 6 before rule number 7, then this would block SMTP access from the external DMZ to the internal DMZ. By placing this rule above however, I can create quite a good rule, in rule number 7 that says I can deliver SMTP mail to the Internet, my internal SMTP server but nowhere else.

This is just one example where rulebase order is important, but this concept applies very much when creating rulebases of this kind that process the rules sequentially.

4. REFERENCES

1. "Private VLANs - A look at Cisco's implementation of Private Virtual LANs (PVLANS)". Cramsession. <http://www.cramsession.com/articles/files/private-vlans--a-look-at-982003-1700.asp>
2. "Cisco Private VLANs". TruSecure/ICSA Labs Firewall-wizards mailing list. <http://honor.trusecure.com/pipermail/firewall-wizards/2000-December/009795.html>
3. "Securing Networks with Private VLANs and VLAN Access Control Lists". Cisco Systems. Updated May 04 2004. http://www.cisco.com/en/US/products/hw/switches/ps700/products_tech_note09186a008013565f.shtml
4. "Private VLANs: Addressing VLAN scalability and security issues in a multi-client environment". IETF. June 17 2004. <http://www1.ietf.org/internet-drafts/draft-sanjib-private-vlan-02.txt>
5. "VLAN Aggregation for Efficient IP Address Allocation" IETF RFC 3069. February 2001. <http://www.ietf.org/rfc/rfc3069.txt>
6. "Private VLAN Catalyst Switch Support Matrix". Cisco Systems. Updated June 18 2004. http://www.cisco.com/en/US/products/hw/switches/ps4324/products_tech_note09186a0080094830.shtml
7. "CCNP BCMSN Exam Certification Guide" Chapter 20. Hucaby, David. Cisco Press. November 2003.
8. "Configuring Private VLANs" – Catalyst 6000 series switches. Cisco Systems. http://www.cisco.com/en/US/products/hw/switches/ps700/products_configuration_guide_chapter09186a008007f4ba.html
9. "Configuring Network Security" – Cisco Catalyst 6500 Series Switches. Cisco Systems. http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a008007e70d.html#wp1043908
10. "Configuring Access Control" – Cisco Catalyst 6500 Series Switches. Cisco Systems. http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a00802011c2.html#wp1119928
11. "Cisco 2600/3600/3700 Series T3/E3 Network Module". Cisco Systems. http://www.cisco.com/en/US/products/hw/modules/ps2797/products_data_sheet09186a008010fba2.html
12. "Cisco 2691 Multiservice Platform". Cisco Systems. <http://www.cisco.com/en/US/products/hw/routers/ps259/ps3145/index.html>
13. "CheckPoint Platform Selection Guide". CheckPoint Software Technologies. http://www.checkpoint.com/products/choice/platforms/platforms_product.html

14. "CheckPoint SecurePlatform – Product Highlights". CheckPoint Software Technologies. <http://www.checkpoint.com/products/secureplatform/>
15. "Cisco VPN 3000 Series Concentrators – Models Comparison". Cisco Systems. http://www.cisco.com/en/US/products/hw/vpndevc/ps2284/prod_models_comparison.html
16. "Cisco VPN 3000 Series Concentrators – Tunnelling and Security". Cisco Systems. http://www.cisco.com/en/US/products/hw/vpndevc/ps2284/products_configuration_guide_chapter09186a00801f1e36.html
17. "Cisco CSS 11500 Series Content Services Switch". Cisco Systems. http://www.cisco.com/en/US/products/hw/contnetw/ps792/products_data_sheet0900aecd800f851e.html
18. "Network Intrusion Detection", Third Edition. Northcutt, Stephen. Novak, Judy. September 2002. New Riders.
19. "ISS RealSecure Network 10/100". Internet Security Systems. http://www.iss.net/products_services/enterprise_protection/rsnetwork/sensor.php
20. "ISS RealSecure Server Sensor". Internet Security Systems. http://www.iss.net/products_services/enterprise_protection/rserver/protector_server.php

© SANS Institute 2004, Author retains full rights.