

Regis University

ePublications at Regis University

Regis University Student Publications
(comprehensive collection)

Regis University Student Publications

Fall 2011

Investigation of Efficient Unified Threat Management in Enterprise Security

Ryan Lynn
Regis University

Follow this and additional works at: <https://epublications.regis.edu/theses>



Part of the [Computer Sciences Commons](#)

Recommended Citation

Lynn, Ryan, "Investigation of Efficient Unified Threat Management in Enterprise Security" (2011). *Regis University Student Publications (comprehensive collection)*. 626.

<https://epublications.regis.edu/theses/626>

This Thesis - Open Access is brought to you for free and open access by the Regis University Student Publications at ePublications at Regis University. It has been accepted for inclusion in Regis University Student Publications (comprehensive collection) by an authorized administrator of ePublications at Regis University. For more information, please contact epublications@regis.edu.

Regis University
College for Professional Studies Graduate Programs
Final Project/Thesis

Disclaimer

Use of the materials available in the Regis University Thesis Collection ("Collection") is limited and restricted to those users who agree to comply with the following terms of use. Regis University reserves the right to deny access to the Collection to any person who violates these terms of use or who seeks to or does alter, avoid or supersede the functional conditions, restrictions and limitations of the Collection.

The site may be used only for lawful purposes. The user is solely responsible for knowing and adhering to any and all applicable laws, rules, and regulations relating or pertaining to use of the Collection.

All content in this Collection is owned by and subject to the exclusive control of Regis University and the authors of the materials. It is available only for research purposes and may not be used in violation of copyright laws or for unlawful purposes. The materials may not be downloaded in whole or in part without permission of the copyright holder or as otherwise authorized in the "fair use" standards of the U.S. copyright laws and regulations.

**INVESTIGATION OF EFFICIENT UNIFIED THREAT MANAGEMENT IN
ENTERPRISE SECURITY**

A THESIS PROJECT

SUBMITTED ON THE 21ST OF SEPTEMBER, 2011

TO THE DEPARTMENT OF INFORMATION SYSTEMS

OF THE SCHOOL OF COMPUTER & INFORMATION SCIENCES

OF REGIS UNIVERSITY


IN PARTIAL FULFILLMENT OF THE REQUIREMENTS OF MASTER OF SCIENCE IN
SYSTEMS ENGINEERING

BY

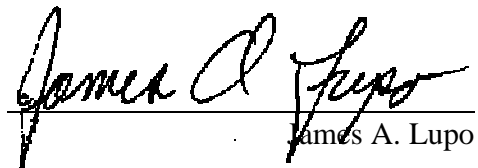


RYAN LYNN

APPROVALS



Paul Vieira, Thesis Advisor



James A. Lupo



Daniel M. Likarish

Abstract

This thesis explores the problems that exist today with perimeter security in data communications specifically the disparate architecture that exists to mitigate risk. Currently there are many different components to the enterprise security perimeter that are not cohesive and do not collaborate well to form an efficient, scalable, operationally supportable gateway design. The thesis breaks down this problem by illustrating the shortcomings of current technologies. These illustrations are used in conjunction with published research and authored research to provide solid footing for the idea of a unified threat management or UTM model. In this model, threat prevention techniques are consolidated into a single logical operating environment that leverages advances in next generation firewalls, intrusion prevention systems, content filtering and antivirus technologies. The results of this investigation are provided in a matrix that shows strengths and weaknesses with a consolidated unified model.

Acknowledgements

I would like to thank God; with Him all things are possible. I would like to thank my wife Dawn who did everything in her power to provide the time, energy and motivation to keep me moving on this endeavor. My children continue to inspire me to keep viewing the world as a child views it – with complete wonder and excitement. Their passion drives me to continue learning. I would also like to thank my friends who have always provided a positive source of input into my life. Their encouragement is what fuels me and they did not disappoint here. Lastly, I would like to thank the faculty and staff at Regis who provided an absolutely wonderful learning environment. Regis faculty has always been attentive and nurturing in my road to obtaining my Masters Degree.

Table of Contents

Abstract.....	ii
Acknowledgements	iii
List of Figures.....	vi
List of Tables	vii
Chapter 1	1
Background	1
Thesis	2
Problem Analysis	2
Purpose of the Study	5
Assumptions, Constraints and Risks.....	6
Chapter 2	8
Introduction to Secondary Research	8
Risk Management	8
Threat Management Models	15
Existing Threat Management Architectures	26
Unified Threat Management	36
Chapter 3	43
Introduction to Methodology	43
Method	44
Design Science.....	44
Evaluation.....	45
Chapter 4	48
Introduction to Results.....	48
Evaluation Matrix	51
Functionality.....	51
Operations.....	66
Support/Cost.....	72
Chapter 5	77
Conclusions.....	77
Chapter 6	81

References.....81

Glossary84

List of Figures

Figure 2.1 Risk Management Process (Kouns & Minoli, 2009)	10
Figure 2.2 Security Systems Lifecycle (Weaver, 2007)	14
Figure 2.3 Onion Model of Security	15
Figure 2.4 Basic Security Model: Routers With NAT	17
Figure 2.5 Basic Security Model: Firewalls	18
Figure 2.6 Basic Security Model: IDS/IPS	19
Figure 2.7 Basic Security Model: Content Filtering	22
Figure 2.8 Traditional Enterprise Security Architecture	23
Figure 2.9 Packet Flow Through Traditional Security Architecture	24
Figure 2.10 Sample Topology: Single Router, Single Firewall	27
Figure 2.11 Sample Topology: Single Router, High Availability Firewalls	28
Figure 2.12 Sample Topology: Active/Active Router and Firewalls	29
Figure 2.13 Sample Topology: Redundant Gateway Designs	31
Figure 2.14 Sample Topology: Full Gateway Deployment	32
Figure 2.15 Packet Flow Through Full Gateway Deployment	34
Figure 2.16 UTM Model: Next Generation FW Core	38
Figure 2.17 UTM Model: Topology of Unified Threat Management	40
Figure 3.1 Lab Topology For Testing	46
Figure 4.1 Magic Quadrant For Enterprise Network Firewalls (Young & Pescatore, 2010)	49

List of Tables

Table 4.2 Unified Threat Management Evaluation Matrix	50
Table 4.3 Functionality – Routing	52
Table 4.4 Functionality – Packet Inspection	54
Table 4.5 Functionality – NAT	55
Table 4.6 Functionality – VPN	55
Table 4.7 Functionality – Voice and Video	58
Table 4.8 Functionality – Content Filtering	59
Table 4.9 Functionality – Antivirus	60
Table 4.10 Functionality – Application Identification	61
Table 4.11 Functionality – IDS/IPS	62
Table 4.12 Functionality – Virtualization	64
Table 4.13 Functionality – High Availability	65
Table 4.14 Functionality – Quality of Service	66
Table 4.15 Operations – Unified Management	67
Table 4.16 Operations – Unified Logging	69
Table 4.17 Operations – Command Line Interface	71
Table 4.18 Operations – Policy Conversion	72
Table 4.19 Support and Cost – Education	74
Table 4.20 Support and Cost – Support	75
Table 4.21 Support and Cost – Cost	76

Chapter 1 – Introduction

Background

Security is the act of eliminating the risk or danger to something. This term defines the ability to protect or keep things safe from harm, whether it takes the human form and protection is offered to people in society or the realms of information security is abstracted and ideas are explored for preventing data breach or loss. Information protection continues to be at a heightened state as organizations continue to spend money to safeguard their core assets (Currier, 2011). This can be most visible in the efforts behind enterprise security architectures which are a sub-set of components that all focus on key areas of the enterprise to offer solutions which mitigate common risk areas. One of the largest parts of the enterprise security architecture is the perimeter defense which consists of both hardware and software tools that provide the fortified boundary of the network. The perimeter is comprised mostly of devices such as firewalls, intrusion detection and prevention systems, anti-virus scanners, content filtering and other mitigation tools.

Largely to this point, many of these technologies have acted in autonomous and specialized ways, focusing specifically on their task. While these technologies perform their assignment, and perform well, much has been said about their relative lack of interaction, synergy and cohesion and how it can actually be quite costly to operate this way (Currier, 2011). Today's organizations are changing from their original landscape to one where volume, both inside traffic and outside of the perimeter continue to grow and application complexity and information security in general become much harder to manage (Cisco Systems, 2009).

Thesis Statement

The focus of this study will be to investigate how the security architecture evolves to meet the demands of modern enterprises utilizing unified threat management in an efficient, scalable and cost effective mechanism. The research will provide valuable insight to enterprises who are interested in the details of unified threat management, illustrate how market leaders are attempting to meet next generation security requirements and advantages and disadvantages of deploying such technologies.

Problem Analysis

In order to really understand why unified threat management is becoming a requirement in today's network perimeter, there is a need to understand what factors in history occurred that lead to this evolution. Early in the Internet's development, academic institutes and research branches of the government, like the Department of Defense and NASA, constructed a web of networks to communicate. Initially it was a risk-free collaboration of groups with a focus on research and learning. In 1988, Robert Tappan Morris, a Cornell University graduate, changed that paradigm by launching the Morris Worm, which attacked NASA and 6,000 other systems (Menninger, n.d.). This event sent shockwaves through the newly created Internet consortium. From this event, network perimeter security was born and the attacks and mitigation techniques would only grow.

The Internet community decided in the early 1990's that having IP routers perform basic access-control was not highly efficient for this function so programming of the autonomous firewall began. The concept of firewalls was introduced with the basic premise of "permitting" or "denying" packets from passing into or out of the network. Although early firewalls were very basic, built for a specific purpose, and not very user friendly, over the years they were tuned to provide more functionality and a better user experience. The first commercial attempt at such a

device with a graphical user interface and mouse came from Check Point Technologies in 1994 with their Firewall-1 product (Check Point Software Technologies, 1994). From this moment to the present, firewall vendors have continued expanding the capabilities of their products to include logging, stateful and deep packet inspection. Stateful inspection means that the firewall is keeping track of each active session and has intelligence into the setup and the teardown of the session. Deep packet inspection allows the firewall to view deep into the payload or data portion of the packet and make decisions on the validity of the packet. Logging also increased the visibility into the firewall by capturing what was being denied or accepted. The firewall has grown up in the past two decades to provide what professionals most commonly think of when perimeter security is mentioned.

The Morris Worm was a wakeup call. Experts realized that one defense mechanism would not be sufficient for every type of security risk they might encounter. As the firewall grew up, so did other security mitigation techniques such as proxies, content filters, intrusion detection and prevention systems and malware or virus detection. Similar to the way that the firewall industry attacked the problem, these other areas of technology followed suit with efforts to make the best possible solution while still remaining largely disconnected from each other. In 1993, Trust Information Systems developed the first application layer proxy which allowed the network to perform acceptance or denial of traffic at the application layer (Cisco Systems, 2009). The proxy has since been extended to meet the demands of thousands of applications and traffic types. When the World Wide Web was constructed, the immediate need to filter the content that users may attempt to reach was realized. An early pioneer in this space was Smartfilter, originally developed by Webster Network Technologies and later bought by Secured Computing, now McAfee. With respect to intrusion detection and prevention, the original Morris Worm

prompted quick development on these systems. Even though the government had been working on intrusion detection software in the 1980's, in 1988 Haystack Labs released the first commercially available product called Stalker (Smaha, 1988). Although a bit immature, this sector would really start to develop when Netranger was released in 1993.

Unlike the previous areas of network perimeter security, anti-virus scanning began on the desktop and transgressed to network scanning appliances. Market leaders began to surface in this space such as Symantec and Blue Coat which offer malware scanning at the perimeter, which alleviates some of the burden off of the end devices.

Many of these mitigation techniques were born and widely developed in parallel with each other but an important note is that most of them were done without much regard for each other. The products remained largely autonomous, with numerous companies each focusing development on their niche. This presents several problems for the sustainability of a security model. The first is that with several points of inspection that a packet must go through, latency and inefficiency will follow. Each of these devices must identify a packet, open it and inspect key aspects of the headers and payload. For each device that was aforementioned, this can mean up to 4-6 devices each slowing the transit of the packet through the network – just to provide security. Couple this with the idea of scalability and bandwidth growth and there is real concern with the enterprise network being able to meet these new demands.

A second issue with this topology is that again since these disciplines were very focused and isolated from each other, there is little or no cohesion with respect to correlating events across the security architecture. If an attack occurs, there is no guarantee that the IDS is able to correlate with the firewall that the event in which they may be flagging, is in fact the same event.

A third concern is operationally supporting this type of network perimeter. Not only are there multiple points of management that need to be accounted for but there are multiple vendors, each with their own management platform. This means many touch points in addition to having staff that is skilled in each one of these areas. Troubleshooting through this type of environment also presents some challenges. Following the packet flow through each device means that many different skill sets must work together to dissect exactly what is happening as the packet traverses the network.

Lastly, the cost model to construct and continue to feed this architecture will become overwhelming. With so many devices handling these functions autonomously and the specialization in the hardware to specific vendors, costs start to become an issue. Each vendor requires hardware, software, maintenance and support. Managing these aspects for one vendor is costly and challenging but doing so for many is not an effective deployment strategy.

From the points listed above, one can discern that the security perimeter architecture has to change. With many suggesting there is a Moore's Law that applies to data traffic rising year over year, a different approach must be taken with respect to network perimeter security (Coughman & Odlyzko, 2001). Unified threat management or UTM refers to the combination of common security procedures into a single and unified system. IDC coined this phrase and it encompasses providing a single pass device that handles firewalling, intrusion, anti-virus, content filtering and other aspects of security disciplines.

Purpose of the Study

As the traffic patterns of network systems continue to change, businesses are put into the position to react and do so quickly to protect their infrastructure. This thesis provides insight into why the technologies available today are not adequate as standalone solutions. Once there is

a clear understanding to what technologies have been available in the past, the thesis outlines new technologies and how they are meeting the demands of tomorrow's networks. Sometimes a generic view of information technologies and security specifically is not nearly enough to provide real value. In that case the thesis evaluates and compares/contrasts some of the market leaders who continue to push the capabilities of security protections. The research uncovers where consolidation of platforms into a common one will increase performance, increase security correlation and reporting all while reducing operating costs.

The results of this research provide interested companies with a current state of capabilities that they may have in use today with added information about other technologies that they may not have investigated in. The research also provides a view of the emerging technologies, what features they bring and what vendors are leading these areas. The information presented can serve as a blueprint for organizations as they move from the current state to an architecture more suited to meet the demands of business with respect to security, compliance and still maintain performance.

Assumptions, Constraints and Risks

The research comes with some assumptions. It assumes that the audience has familiarity with some or all of the different types of threat management. As stated above, many of these technologies are installed into the enterprise as standalone devices. In order for the research to offer positive value, the audience should be familiar with these mitigation techniques and will ideally share the same opinions of the problem analysis, that existing technologies are not well positioned to be successful.

There is an assumption that the testing in the results section of the thesis is a snapshot of the vendors' capabilities in time and under nominal conditions. Vendor technology can change

frequently to improve numbers and traffic patterns can alter results. The results are intended to give a general overview of the platforms and more-so provide a common trend with respect to where the architecture is going.

The research does not come without possible constraints. An obvious problem is that one size does not fit all when it comes to enterprise security. Size, complexity and other characteristics of data can all alter the needs of a company. Based on the enterprise need, one area of unified threat management may hold more value than another. An example here might be the need for web content filtering. This is merely scanning the traffic that is destined for the world wide web and ensuring that based on pre-defined categories of acceptable traffic, users are not accessing content that are against company policies. In the case of a large enterprise company, traffic demands for throughput may be much higher than a small business. In this case, consolidation of this function into a unified solution may not be able to scale well for the enterprise company.

Another constraint would be legal requirements that each company may have to adhere to. Things such as PCI compliance, HIPAA and SOX may also shape the needs of a company. Another area where special concern needs to be addressed is in government systems. The government is bound by their own set of special rules for classification, protection and securing of data. For the purposes of this research, the evaluation is not bound by any of these.

Chapter 2 – Secondary Research

Introduction to Secondary Research

Getting to the heart of threat management, it is important to take a step back and investigate the root of risk and how it is managed. Risk is “the potential that a chosen action or activity (including the choice of inaction) will lead to a loss (an undesirable outcome)” (risk, 2011). Risk can be present in human behavior or within decisions that guide a business to become vulnerable in some capacity. The truth is that risk has several different contexts depending on which facet of business or what discipline it’s being represented by. With respect to this thesis, risk is any potential unwarranted or undesirable interaction with enterprise data.

Within this area, a set of policies are created that outline potential risks and attempt to quantify how an organization may avoid or react to such activities. The subsequent sections show how risk management ultimately creates policies to handle threats and how the types of threats can be categorized by different parts of the enterprise architecture. For the purposes of this research, the interest is in the network perimeter, or the outer defense architecture and the manner at which it has evolved to the current state. More-so the research shows how risk management is struggling to meet current and future threats with the existing perimeter technologies.

Risk Management

According to Weaver (2007) risk management is a term to describe the process of identifying, choosing and setting up countermeasures justified by the risk identified. One of the important things to take away from this meaning is the word “process”. Since risk management in the context of an organization is the focal point, there needs to be some type of process that

makes this whole thing function. Without a formal process in place, management cannot take place and more importantly risk will not be identified.

Risk management is not a new term. Looking back over history, it could be argued that risk management existed as long as humans made provisions to deal with a potentially bad situation. In the 1700's, in ancient Babylon, risk management was exemplified in pre-paid loans that merchants would secure in order to insure the transportation of goods over long distances (Hubbard, 2009). This early form of insurance was a very primitive form of risk management. From the 1700's to now, risk management has largely existed in the financial, insurance and government sectors. No matter the application, the process to ensure a stable "norm" has been recognized as a much needed process. Much of these earlier examples of risk management were specific and not standard across different applications.

By the 20th century, the international standards organization or ISO began to see a uniform need across common businesses for some type of management of this risk. There are several ISO standards that document systematic approaches to risk management across many different types of businesses. If you analyze these methods, they essentially have the same steps. In figure 2.1, a general view is given of a risk management process that allows us to understand how the complexities with risk management are quantified (Kouns & Minoli, 2009).

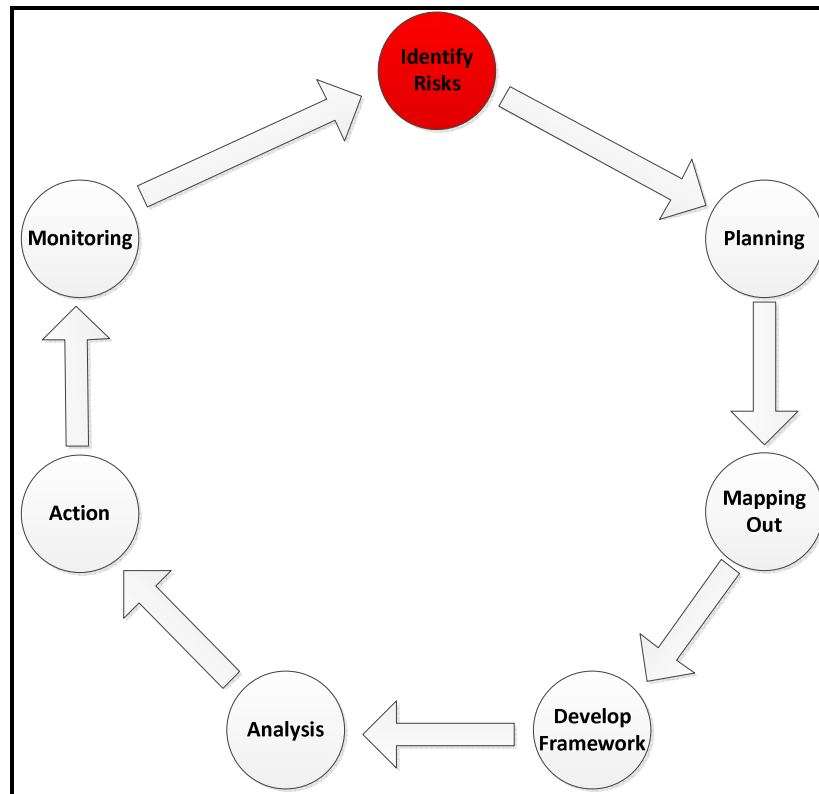


Figure 2.1 Risk Management Process (Kouns & Minoli, 2009)

In revisiting the definition of risk management, the word “identify” is used twice in the sentence. This is the most important step of the risk management process. Identification is the awareness of the risk. Without awareness, reaction and mitigation are challenging. Charette (1996), in his paper about the importance of identification in risk management, points out that without solid footing on what the risk is, misidentification can lead to not only missing the obvious threat but also investing a lot of resources into a misidentified risk. He strengthens this with an example from the health care industry which extrapolates to show that many life saving pharmaceuticals have been withheld from the market for misidentification of risk while many potentially deadly ones have been distributed to the masses. This error in recognizing risk has the potential to undermine the entire risk management process as interesting but possibly irrelevant.

As it applies to enterprise security methods, risk management has struggled to keep up with the types of threats that are being released at a constant rate. Data communications has made it easier for attackers to thwart holes in the perimeter and have kept security professionals reactionary. With so many different types of applications that the enterprise network must transmit data for, technologies have been created to protect any unwanted access. At the core, risk management succeeds only if step one of proper identification of the risk is accomplished.

The planning phase of the risk management cycle refers to the outline for how the process will flow for the remainder of the exercise (Kouns & Minoli, 2009). This would involve what type of high level method will be used to gather, evaluate and assess the risk. It is the part that is quite unique to each company because in the planning phase, may be elements that are specific to a business sector or type of organization. For example, a government agency that may be looking to utilize this cycle may have its own set of processes and procedures that must be adhered to that would be identified during the planning phase. Essentially planning involves taking into account the business environment and possibly already established processes for the execution of the remainder of the cycle.

Mapping out the risks involves a few different steps. The first is to identify who the stakeholders are that have vested interest in the risk (Kouns & Minoli, 2009). It cannot be up to the risk management group to decide what priority is put onto a risk. Stakeholders need to show business reason and potential damage that an identified risk could have. It is also in this step that the criteria for how risks will be interpreted should be outlined. Each evaluation needs to be grounded by a common perspective so that an apple to apple comparison can be done.

Defining the framework allows a systematic approach to be used to evaluate the risk and ensure that entire process is handled in the same way (Kouns & Minoli, 2009). This ensures

consistency when analysis is done. There are several different frameworks that can be used that range from standardized by international bodies to ones that are homegrown and customized to a specific business type.

The analysis of the risk is also a vital step. Risk management is responsible for putting into place policies, mitigation techniques and technologies to minimize or eliminate the risk. Without performing this level of analysis, with the above mentioned stakeholders, risk management is put into the position of guessing which ones they believe are critical. Without having a deep understanding of the business, its processes and how the operations work, risk management could be entirely off base.

The action part of the risk management process is the mitigation or solution to the risk. There are several courses that an enterprise can explore with respect to this step. The first and obvious step is avoidance of the risk (Hubbard, 2009). This is to say that the company decides to not put themselves into a position of risk in the first place. Maybe this means they do not release a certain application. It could also mean that they decide not to allow a certain type of traffic. Transferring the risk involves using another external source to carry the risk. An example of this might be in the case of Payment Card Industry or PCI compliance that gets outsourced to another company. Transferring the processing of credit card data to an external company places the risk on that company. Transferring is not always a good idea though if you consider how important information is to an enterprise. This prized asset leaving the corporate walls can often be a difficult decision to make. Another less active solution to mitigation is to simply accept the risk and do nothing (Hubbard, 2009). This decision should be made with a thorough analysis of the risk. If it is identified that the risk carries low probability and low impact, then it may be in the best interest of the company to document the risk but ultimately accept it.

The most popular choice for mitigation though is to reduce the risk. This is where threat management and network perimeter security attempt to provide the organization with a reduction in risk. Because systems are not perfect and only perform within specific rules, they can be compromised, overcome and circumvented. Threat management attempts to take the known risk and with analysis of the probabilities, place technologies into the best parts of the network to reduce overall risk. The thesis speaks to the weakness in technologies today in how they are deployed autonomously which creates operational overhead, complexity, scattered view of the architecture and problems with correlation of events. The next section will discuss in detail how each area of the network perimeter for threat management evolved to its current state and where they have failed to provide enough value.

The final steps of the risk management process involve actually implementing the technologies and more importantly monitoring the solutions for effectiveness. It is not enough to identify, analyze and implement only to walk away with a false sense of security. Monitoring of the technology is the quality assurance that the security industry so desperately needs. Monitoring provides the risk management team with feedback about what is working and what is not. Because of the importance of this step, the process for risk analysis is iterative. Figure 2.2 shows that once the system is activated, new risks would travel through a cyclical process that forms the operational model of the enterprise security architecture. As is drawn out, the cycle from figure 2.1 is represented in figure 2.2 as the “risk analysis” which drives the actual security policies that will be incorporated into the overall architecture (Weaver, 2007).

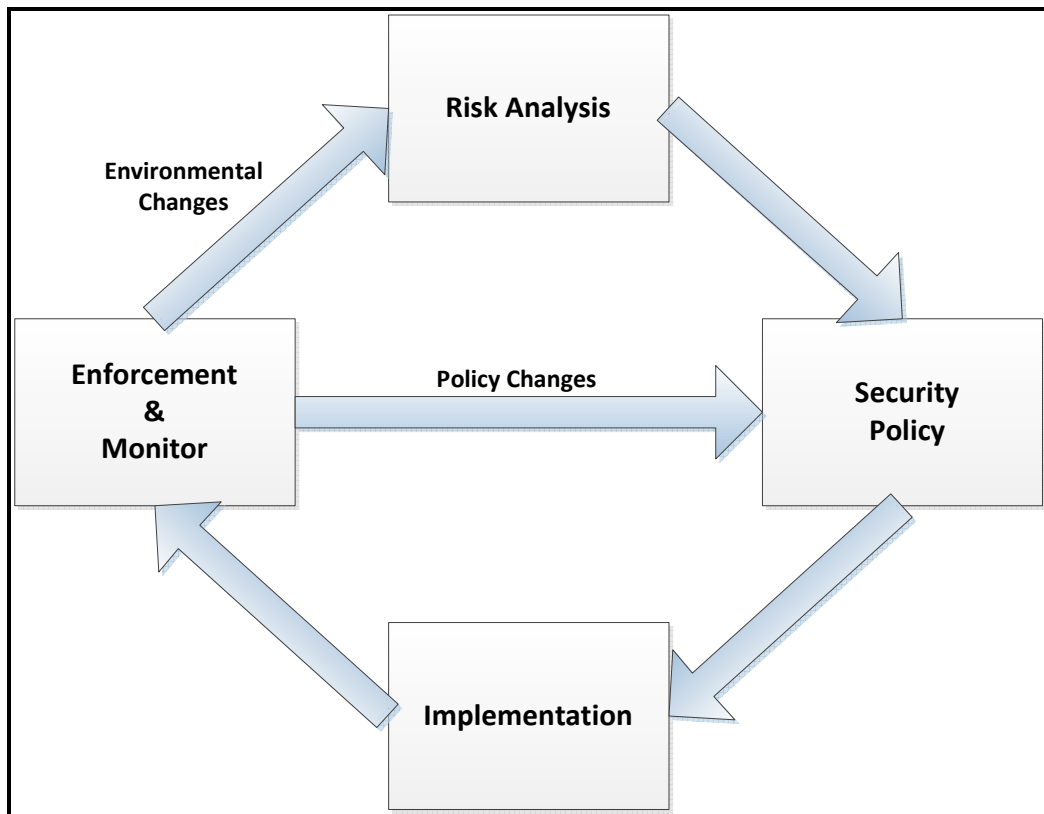


Figure 2.2 Security Systems Lifecycle (Weaver, 2007)

These policies could be rules and acceptable uses of company resources but more importantly outlines what technologies will be implemented to handle the threats that were identified. As mentioned above, after implementation is complete, enforcement of the policy will result in reporting that provides indications of where the implementations were successful and where refinement is needed. Refinement is fed back into the security policy while any changes to the environment are pushed back to the front door of the process.

Because the feedback of the security technologies is so powerful, the research will show in later sections how reporting of the data, which can be massive depending on the company, is critical. In many enterprise networks, the reporting is not centralized, not analyzed and in some cases is never looked at.

Risk managers however do not have an easy job by simply focusing on a single part of the enterprise architecture. There are several sections of the puzzle that each needs to circle the process outlined in figure 2.2. These threat management models each have their own set of challenges but share an overlap as they depend on common technologies to solve their respective problems.

Threat Management Models

Threat management is a derivative of risk management. In order to be clear and concise about each of their meanings, the difference between the two should be clarified. A threat is anything that can exploit vulnerabilities and obtain, damage or destroy an asset. In this case the asset is information. Risk is the probability that a threat will exploit these vulnerabilities. So you can see that in order to effectively manage risk, there must be evaluation of the threats that our architecture faces. Threat management in the realm of enterprise data architectures can be subdivided into an “onion” diagram. In looking at figure 2.3, it is highlighted that at the heart of the model is the data itself. This is our core asset. As the onion is peeled back, the data interacts with applications that reside on hosts which ultimately can send the data to another host by using the network.




Figure 2.3 Onion Model of Security

Each layer of this model has risk and ultimately will have vulnerabilities that can be exploited by threats. Threats can be an application flaw that allows an attacker to gain access. It can be a hole in the operating system that can be compromised. For the purposes of this thesis, the focus will be on the network layer which is commonly called the “network perimeter”. The network perimeter model has changed quite a bit over the years to coincide with the rise of global business and the exchange of data. The initial protection against threats was mostly a lack of options. When systems were not connected to each other, the system was considered closed and thus the threat level was low. With the explosive growth in businesses exchanging their data, evolution of protecting the closed system occurred.

The router is the first layer of this defense (Al-Radhi, 2009). A router is simply a device that receives and sends data packets to and from a source and a destination at the network layer. In the perimeter, a router is commonly used to connect the “trusted” enterprise network with an “untrusted” external network. Initially when these devices were used to interconnect networks, security was primitive and came in the form of an internal firewall to protect. Routers however have evolved to the first point of security protection for a company.

Routers have a few functions that they specialize in. Obviously routing is a key component to moving data from point A to B. If routing was compromised, traffic would not be able to transit so protecting the process that routers were designed for is paramount. Routers are not impervious to vulnerabilities, so protecting the routing infrastructure is part of the overall security architecture.

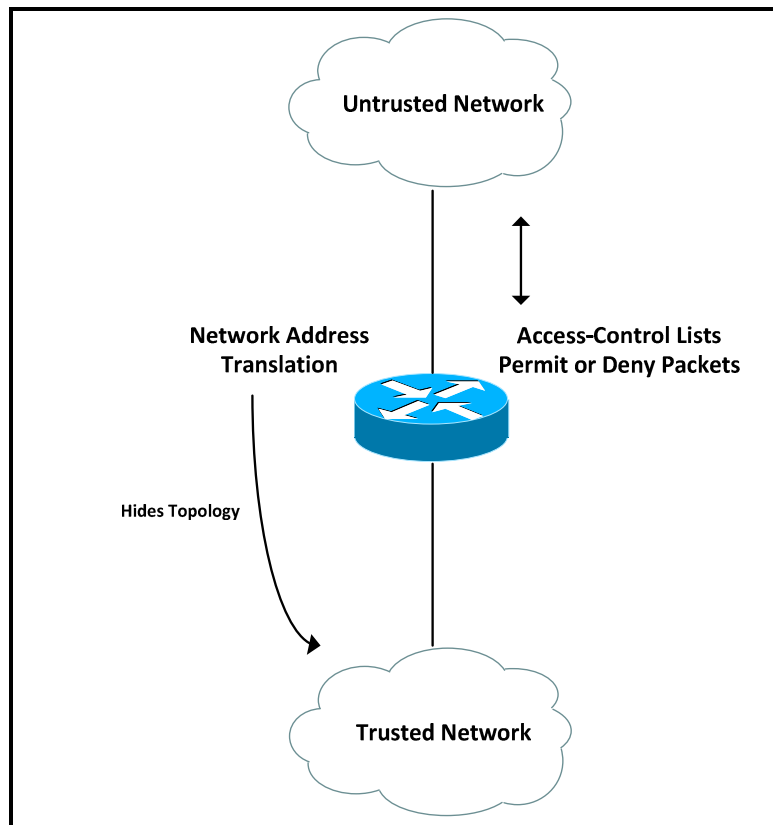


Figure 2.4 Basic Security Model: Routers With NAT

Routers were also originally used to perform functions that are known today to be implemented by firewalls. Access-control lists are basic firewall rules that allow the router to permit or deny traffic based on IP addressing and layer 4 ports (Al-Radhi, 2009). Couple this with the router's ability to perform network address translation or NAT and companies had very primitive forms of the commonly viewed firewall today. Figure 2.4 shows these basic functions. As external connectivity continued to grow, the router became more cumbersome to configure for protection and to perform the NAT functions.

Firewalls were created to relieve the routing platforms of this burden. They were more purpose built to handle controlling access to and from the company network.

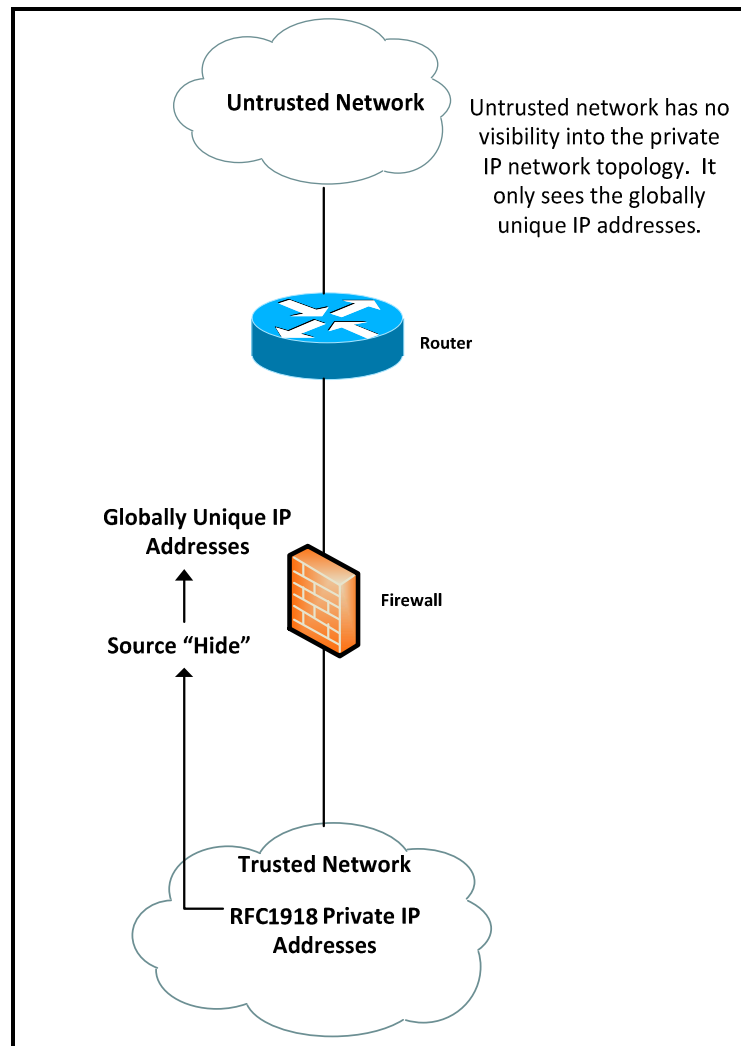


Figure 2.5 Basic Security Model: Firewalls (Lynn, 2011)

Their focus was on making it easy to administer rules that with the course of technological history were starting to get quite complex (Forrest & Ingham, 2002). They further protected the enterprise by providing the basic need for NAT which was discovered to be a very valuable security tool in itself. Originally NAT was designed to connect a corporation's private internal network to a public network such as the Internet. In order to do this, the addresses needed to be translated. In translation, the entire internal private network is masked from the outside public network, thus providing a sense of "hiding" the topology as shown in figure 2.5.

While firewalls originally were excellent at filtering packets with basic rules and criteria, they were still vulnerable to someone spoofing traffic to appear to be legitimate (Forrest & Ingham, 2002).

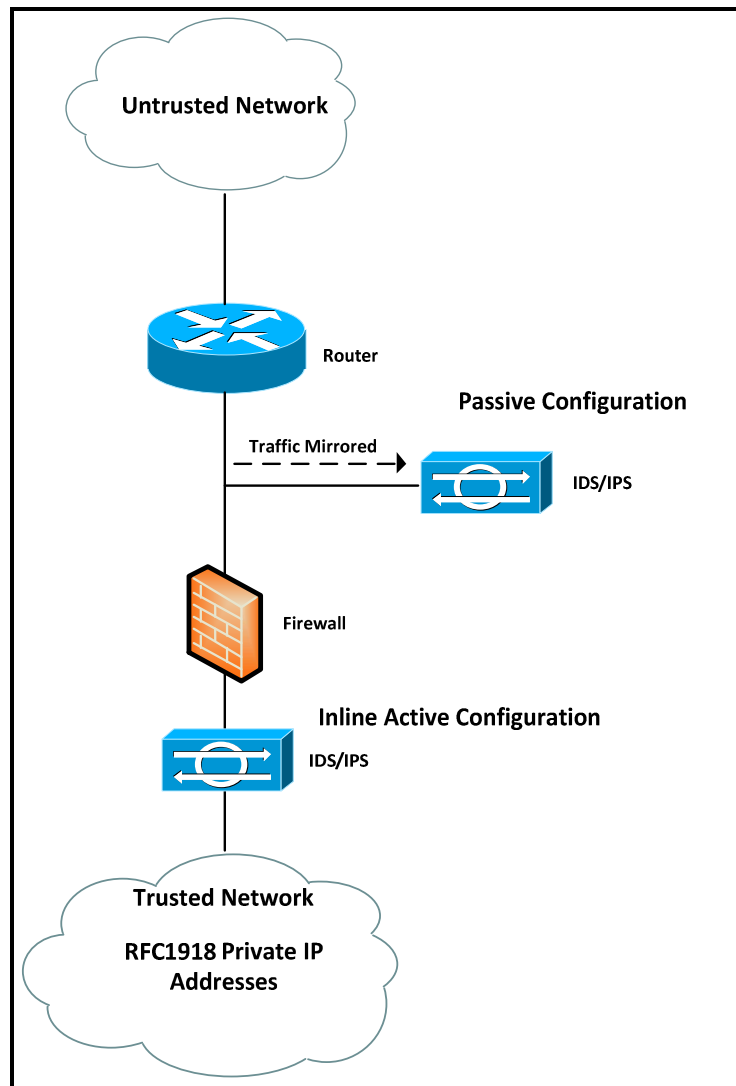


Figure 2.6 Basic Security Model: IDS/IPS (Lynn, 2011)

With the rise in spoofing attempts, firewalls evolved into more intelligent devices by tracking the sessions that are set up when data communication occurs. This involved watching the traffic and paying attention to the setup and the tear down of the session. If the firewall sensed that the traffic had been manipulated, it could then react to it.

Although the firewall was well positioned to scale to handle new types of traffic, the hardware advancements for new features and functionalities were lagging which caused more purpose built platforms to spawn. Intrusion detection systems or IDS attempt to identify traffic that is intended to breach the integrity of the system. They watch the network streams and look for intrusions that are not authorized. They do this in two basic ways. The first way is through traffic signatures which are copies of what the attack looks like that are stored in a database on the IDS. While the IDS is watching traffic, it compares the traffic patterns to this frequently updated database. If the traffic matches the signature, notifications and alarms are sent.

The second type of IDS is one of mathematical anomaly detection. The IDS is instructed to build a baseline of what “normal” traffic patterns look like. It uses this information and statistics to find deviations from the norm. Once it is detected, the IDS can notify that a compromise is in progress (Innella, 2001). An IDS is normally not intrusive and does not become an intrusion prevention system or IPS until it proactively takes action on the attack. IPS refers the system’s ability to react and defend the network by denying the traffic from passing through.

Also different from a firewall, an IDS can also be placed internally inside of the perimeter in strategic locations to identify internal intrusion. Figure 2.6 shows how the IDS/IPS platform operates within the external gateway environment or inside of the network. It also shows that an IDS/IPS can operate by simply monitoring traffic and without being in the actual traffic flow. When an IDS/IPS is put “inline” with the traffic, all traffic is flowing through the device and this can introduce another layer of failure into the network (Innella, 2001).

Similar to the way that IDS/IPS evolved from a specific need that firewalls could not fulfill, the antivirus/antimalware devices were created to contend with the large upswing in

viruses that were finding their way into the enterprise network (Doctor & Poynter, 2003).

Viruses embed themselves into the payload or data portion of the packet and firewalls were not well equipped with the hardware needed to process looking that deep into the packet. Antivirus appliances were engineered and deployed to be collocated with common applications that house viruses such as email or web browsing. Email messages are common places where viruses are introduced and early technologies were not positioned to catch these before users opened the attachments, releasing a virus onto the internal architecture. Once released, these viruses can steal corporate data, open holes for remote access and also cause denial of service attacks that could render the network unusable. The topology for antivirus appliances is very similar to IDS/IPS solutions represented in figure 2.6.

Content filtering was also developed during the same time in order to provide a level of restriction to what websites a user could access. Content filtering is a database of sites that are denied or blacklisted which are filtered to prevent users from reaching those sites (Doctor & Poynter, 2003). The databases are continually updated as new Internet web sites are created. Content filtering traditionally has two methods of deployment, similar to the way that IDS/IPS and antivirus appliances are deployed. Both are shown in figure 2.7. The filters can be in-line of the flow of traffic which can be more intrusive if there is a failure on the appliance or they can be deployed in a redirected fashion where traffic is matched and then redirected. Unlike IDS/IPS, content filtering redirection has a level of failover in that if the content filter is in a transparent mode and fails, the traffic can fail straight through. This failure would put the topology in a scenario where there would be no protection during the outage.

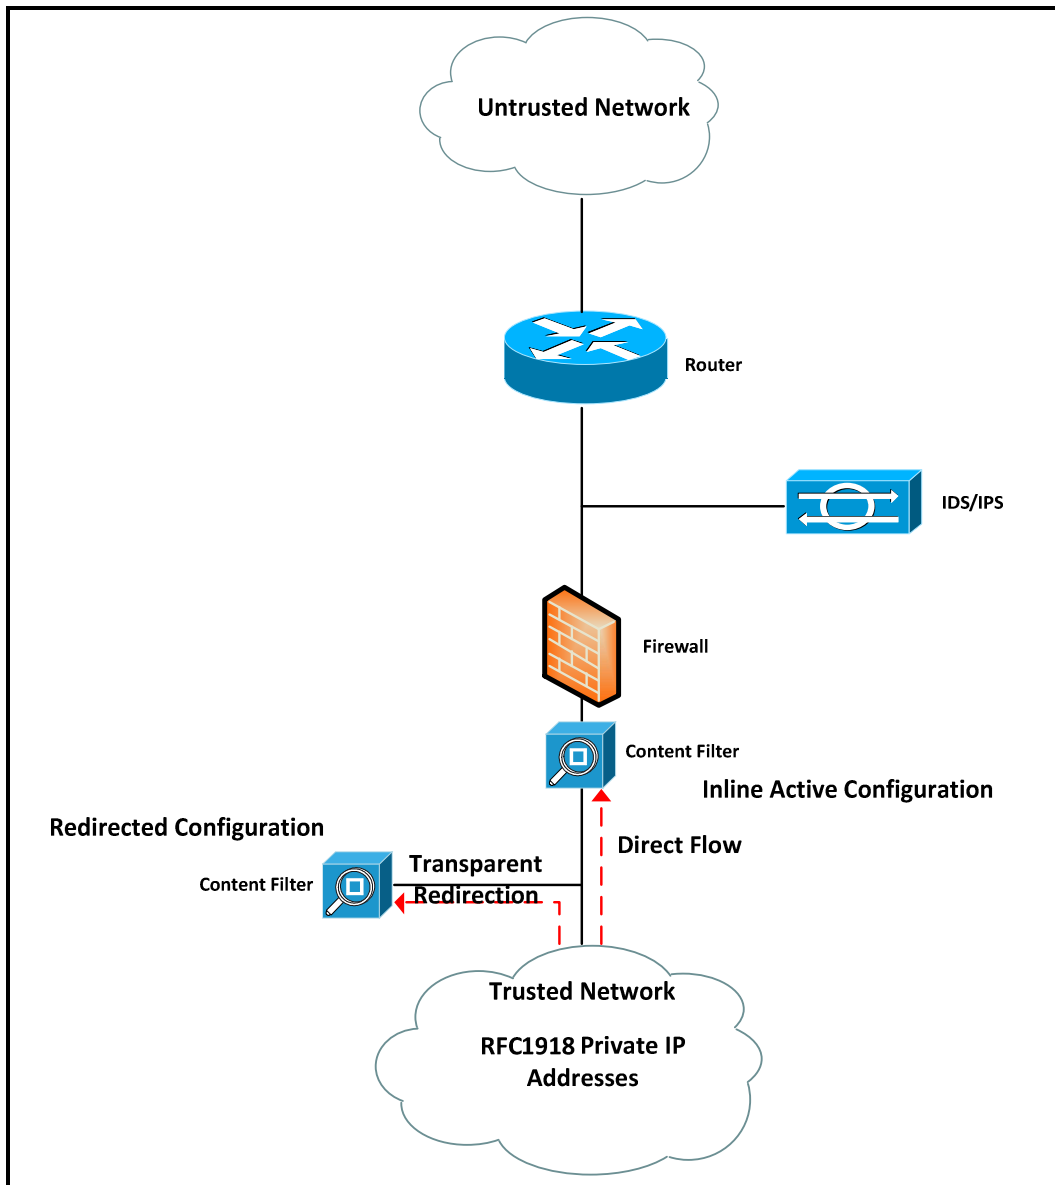


Figure 2.7 Basic Security Model: Content Filtering

The aforementioned technologies are the building blocks of the common enterprise security architecture as depicted in figure 2.8. The size of the corporation can increase or decrease the scale of these devices depending on need. The takeaway from this section is that there is no single platform to handle the various types of risks that existed. Businesses had to

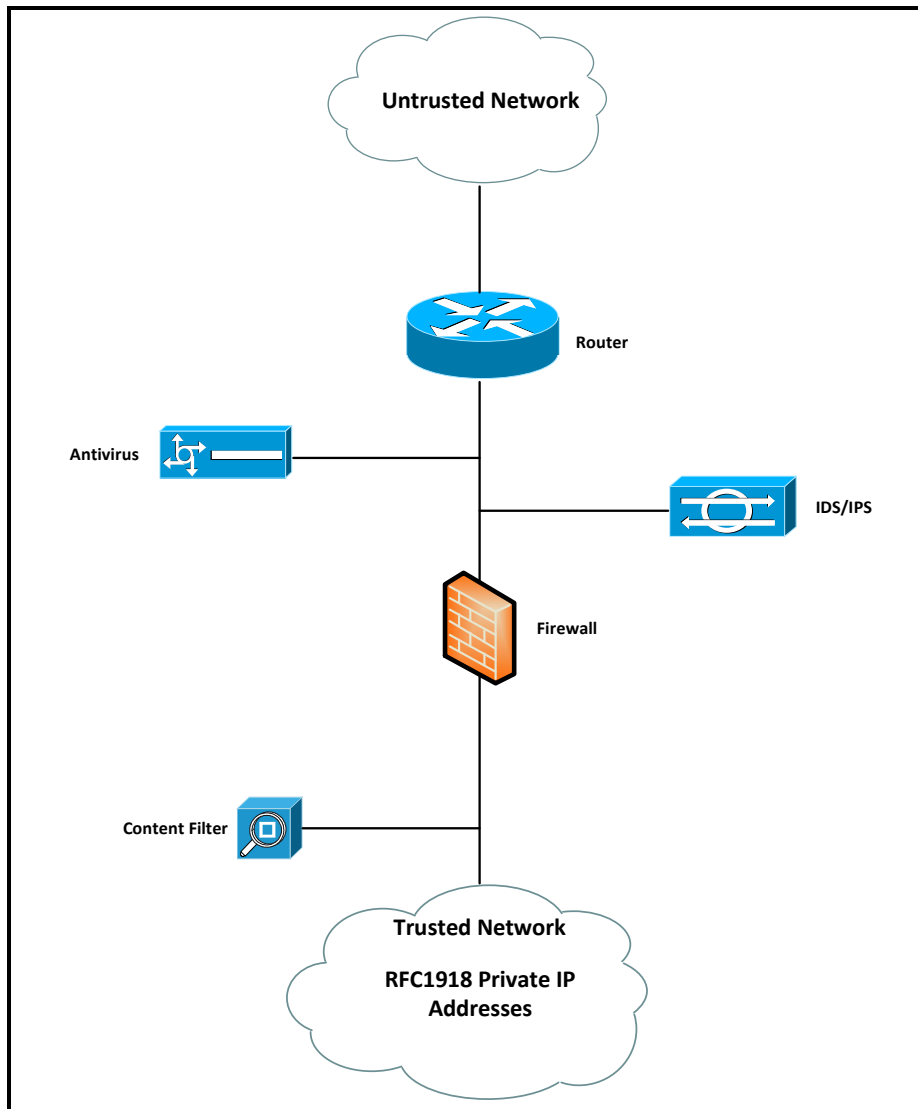


Figure 2.8 Traditional Enterprise Security Architecture

rely on different hardware, software, operations, maintenance, security policy, reporting and other elements of an autonomous system. If this wasn't burdensome enough, performance through this type of network suffered multiple inspection points.

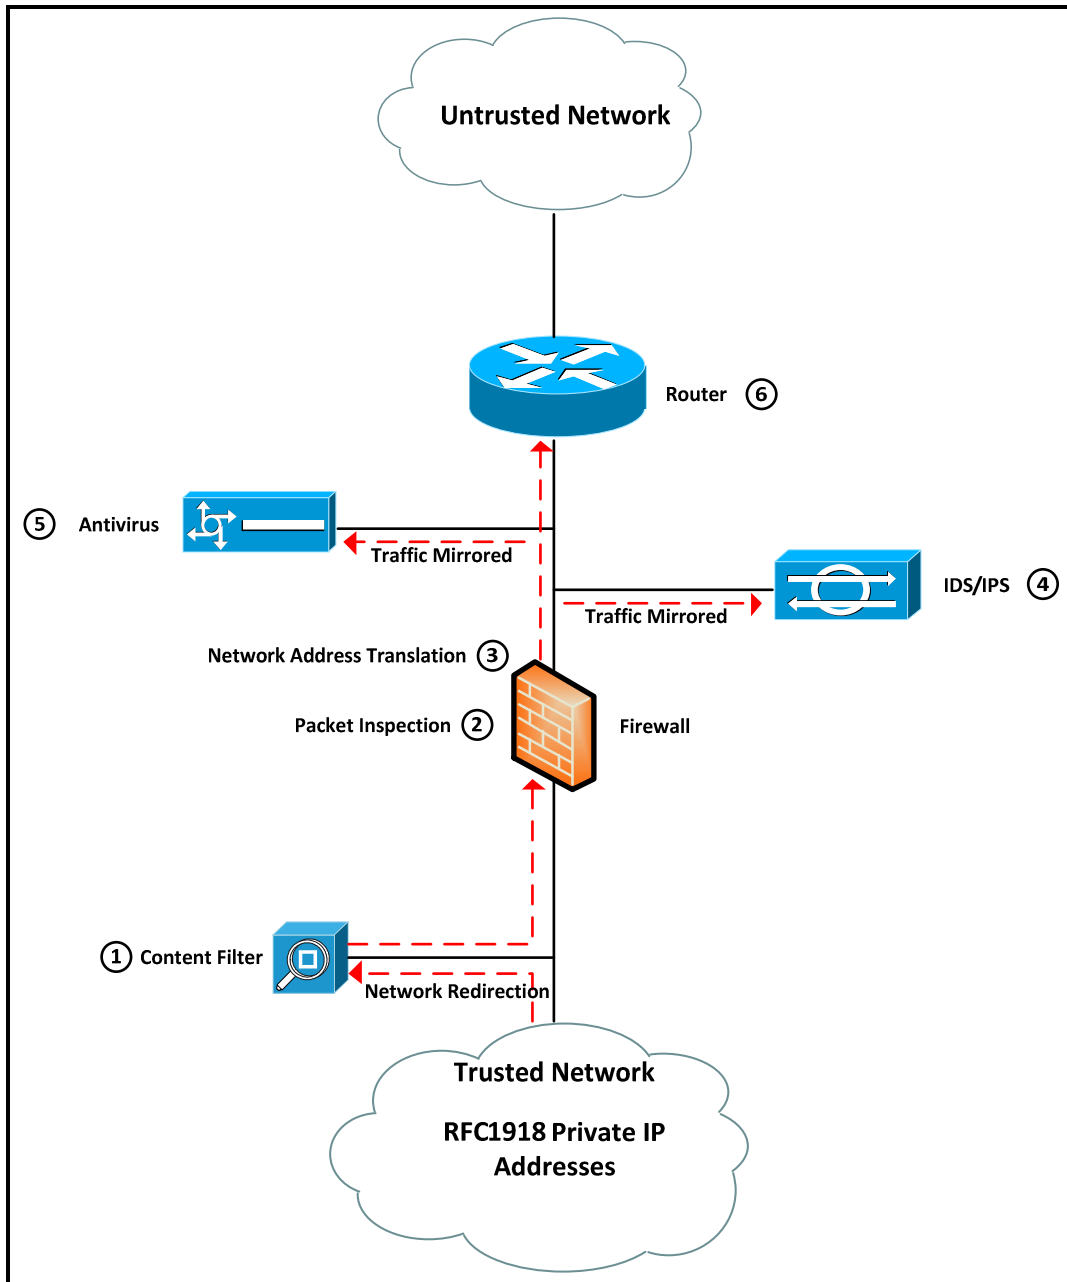


Figure 2.9 Packet Flow Through Traditional Security Architecture

Figure 2.9 shows the typical traffic flow of a single packet in which at every stop, the packet is delayed by being opened and evaluated on the criteria that the specific device is responsible for. In device number 1, the packet is received by the content filtering device, which opens up the packet to evaluate the web site being requested. Once it completes its decision, the

content filter will reassemble the packet and send it device number 2 which is the firewall. The firewall, depending on how sophisticated it is, will open the packet up to at least identify the source, destination and ports for the communication session. Some firewalls will look deeper into the packet for some rudimentary defenses at more intelligent layers. Once the packet has been evaluated, it then has the option to be translated by the firewall. Step number 3 could change the source, destination or port numbers of the communication.

The IDS/IPS device now watches the traffic leave the firewall. This could be in a passive configuration where it is not in the active path of the data or in an active configuration. The same is true of the antivirus device which is attempting to inspect for viruses or malware. Finally the router in step 6 is able to receive the packet and make a determination of where to route the packet. As packets flow in the opposite direction, the same devices are evaluating the flow. Efficiency of expediting the forwarding of the packet comes into question here. With many devices in the flow of the traffic, the potential for opening the packet numerous times can be quite high which will introduce latency along the way.

Couple this with the operational headache of supporting multiple vendors each with their own platform. Experts in each appliance would need to be kept on staff to support the individuality of the solution. Operational complexity would also increase as the packet is redirected to each appliance. Experts would need to understand how each device ingests information and exports it back onto the network. Network analysis becomes vital at this point to be able to determine how the traffic should flow. Correlation of the individual products also could result in a manual task which could be different across platforms. Timestamps are largely relied on today as the only form of correlation.

If multiple vendors are used in the security perimeter, individual support contracts would need to be in place to support each vendor's hardware and software. This also means a different process to follow in handling outages or incidences depending on the number of vendors.

Space, power, cooling and cabling to each device may seem like a small hurdle to overcome but for large implementations or ones where space is a premium, this can increase the operational costs as well.

Many of these issues plagued the industry for many years and were only exacerbated by the sheer increase of attackers taking advantage of the industry's scattered approach to perimeter security. Meanwhile the market for hardware based accelerated services started to catch up and what was originally looked at from the firewall started to become a second attempt at consolidating some or all of these platforms. Unified threat management is the realized ability of the firewall to evolve into a next generation platform that is capable of performing deep packet inspection, NAT, intrusion detection and prevention, anti-virus, anti-malware and content filtering.

Existing Threat Management Architectures

In order to understand the need for a unified threat management model, there needs to be some analysis why the aforementioned technologies have limitations in scalability, cost, management, operational support and efficiency of the system itself. For this purpose a base level design will be used to show how as the traffic or requirements for security increase, the limitations above will surface.

Figure 2.10 shows a basic security model that includes a router and a firewall that protects the trusted side of the network from the untrusted side.

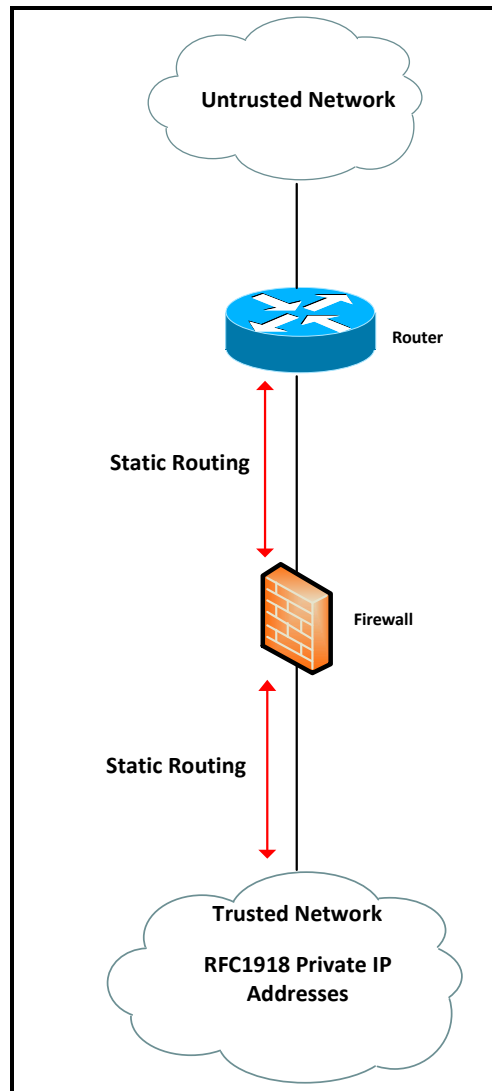


Figure 2.10 Sample Topology: Single Router, Single Firewall

In this model, the complexity remains fairly low. The router is performing its obvious function of routing packets in and out of the environment but also provides a rudimentary first layer of defense with access-control lists that stop unwanted traffic before it even gets to the firewall. Because of the simplicity of the topology, there is no dynamic routing between the firewall and router, which provides very low operational support as far as complexity is concerned. This design may work well for businesses that do not have stringent requirements for redundancy or other forms of security.

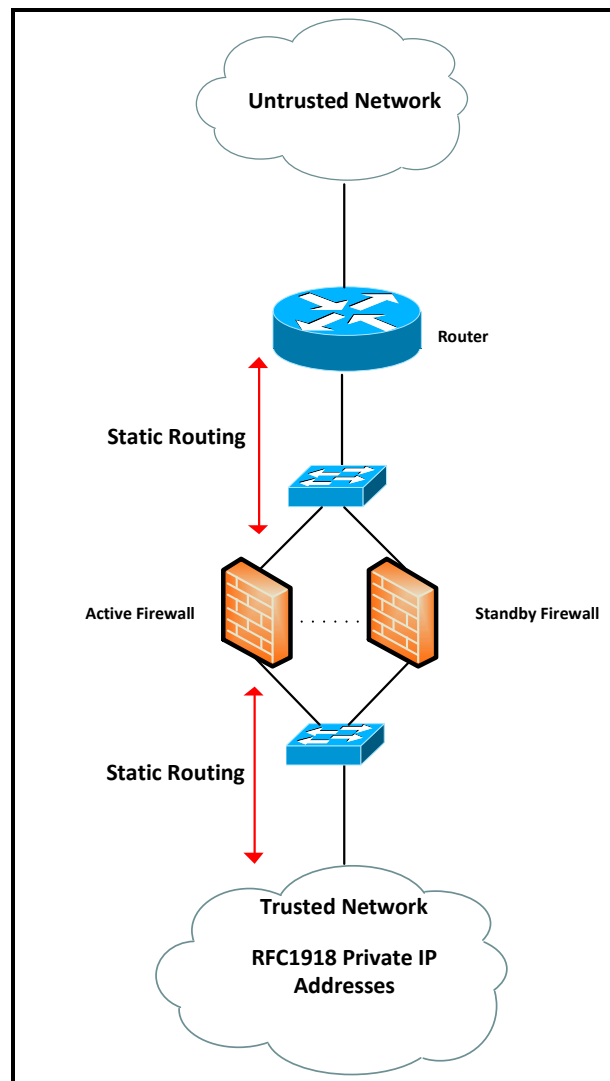


Figure 2.11 Sample Topology: Single Router, High Availability Firewalls

Figure 2.11 shows the growth of this topology when redundancy is a requirement. Bandwidth needs still dictate at this point that a single "active" firewall is sufficient but in the event of a failure, there is a requirement to have a standby firewall that can take over as master. From figure 2.11, the observation is that the environment is growing in operational support, now with two firewalls operating in a high availability or HA configuration. In this design, the option is still open to keep the network flow simple by not invoking any routing protocols on the

firewalls themselves as in the HA configuration there is still only a single active firewall at any given time.

If bandwidth demands increase to utilize the capacity of a complete firewall, the topology must change to accommodate this increase in bandwidth. Figure 2.12 has now been replicated to handle the bandwidth needs but it also may have the need to invoke some routing awareness of both sets of firewalls.

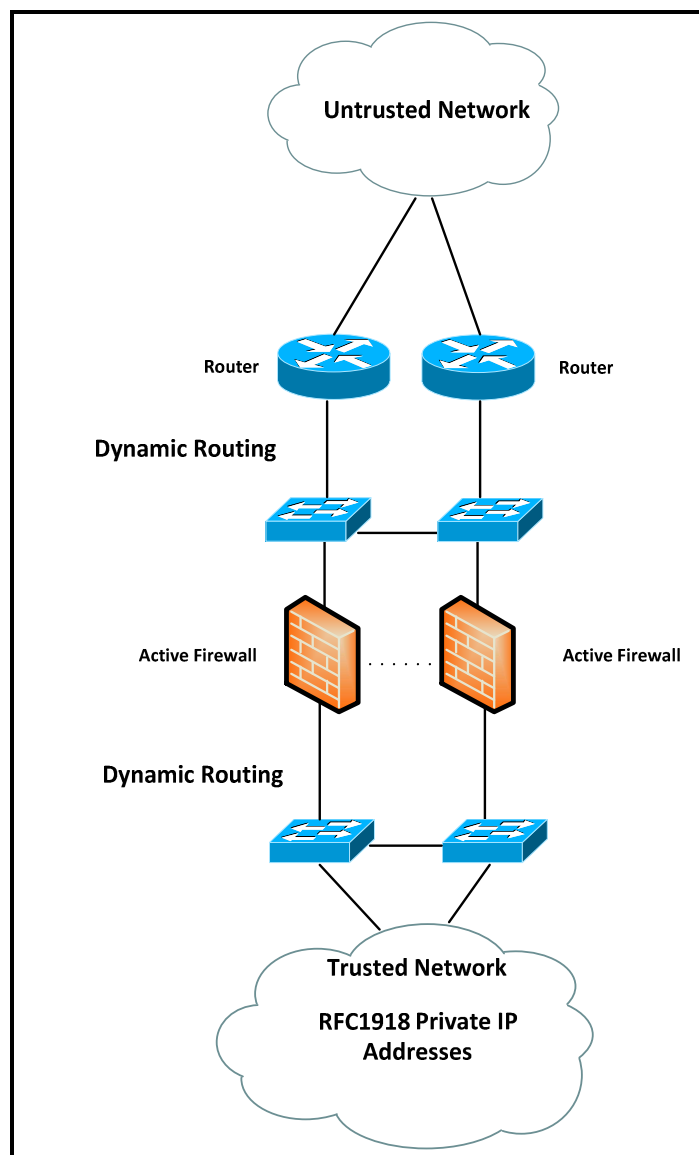


Figure 2.12 Sample Topology: Active/Active Router and Firewalls

In the diagram, both firewalls are actively sending and receiving traffic. This may require the firewalls to start routing dynamically through protocols that are industry standards such as OSPF, RIP or BGP. Operational complexity is beginning to rise now with two routers and firewalls both active in routing. An obvious cost increase will occur with any additional equipment that is added. This is both from an initial capital investment and a continuing operating expense with yearly maintenance and support. As bandwidth continues to rise, which is substantiated by Moore's Law, the scalability becomes costly and companies are forced to optimize what they can with what they have (Coffman & Odlyzko, 2001).

In figure 2.12, both firewalls are active but this may leave some risk in that with both firewalls fully loaded with traffic during normal operating load, that a failure of one of the firewalls would overwhelm the non-failed firewall. Because of this, in order to scale the topology but still provide an active/standby scenario, the network would need to operate with two replicated silos, with the traffic split between the two. This would also be the case if there were different untrusted networks that the corporation needed access to. For example, many companies have access to the Internet in addition to access to a third party vendor that the company does business with. In that case, there could be a need to separate the traffic. Figure 2.13 shows the topology with this type of need. Obviously, the equipment costs are apparent in that the silo has been replicated twice. Operationally, the staff is now responsible for double the equipment which begins to introduce complexity into the environment.

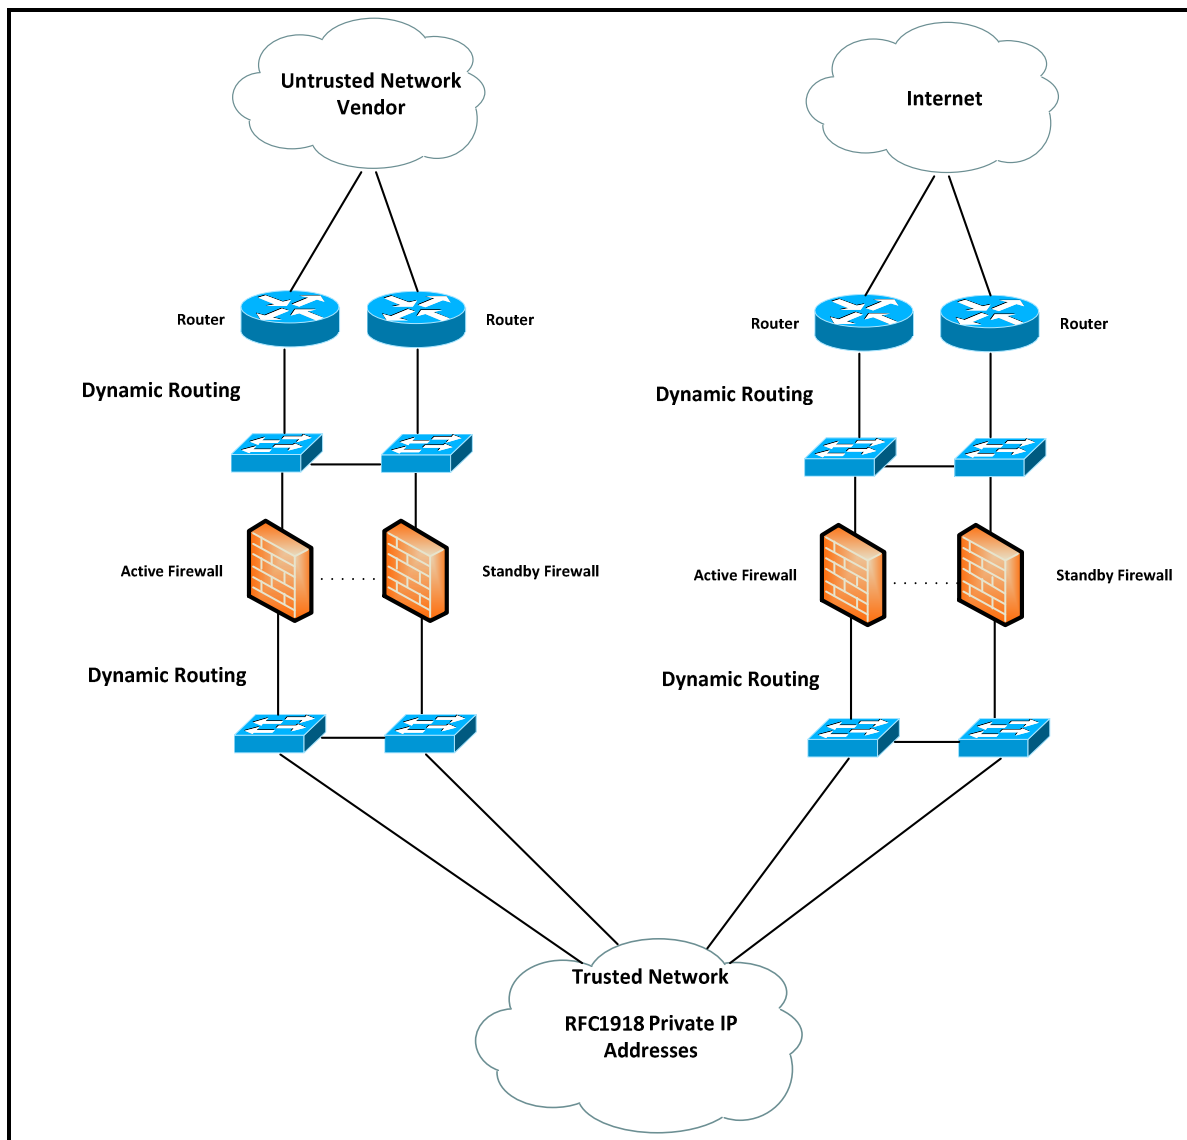


Figure 2.13 Sample Topology: Redundant Gateway Designs

The research so far has outlined merely the basics of a firewalled environment. Because security is not a one size fits all topology, the architecture must expand to combat threats that come in the form of intrusion, viruses, malware and the filtering of content that users inside of the network perimeter are viewing. Geo-redundant deployments to account for disaster recovery should also be looked at. Assuming the technologies could be added one at a time, skipping to the full scale deployment demonstrates the size and breadth of what these security tools demand

on the overall architecture. Figure 2.14 depicts a two gateway design with redundant highly available firewalls, routers, intrusion prevention technology, anti-virus protection and content filtering for outbound traffic. The obvious observation is that the device count has risen significantly.

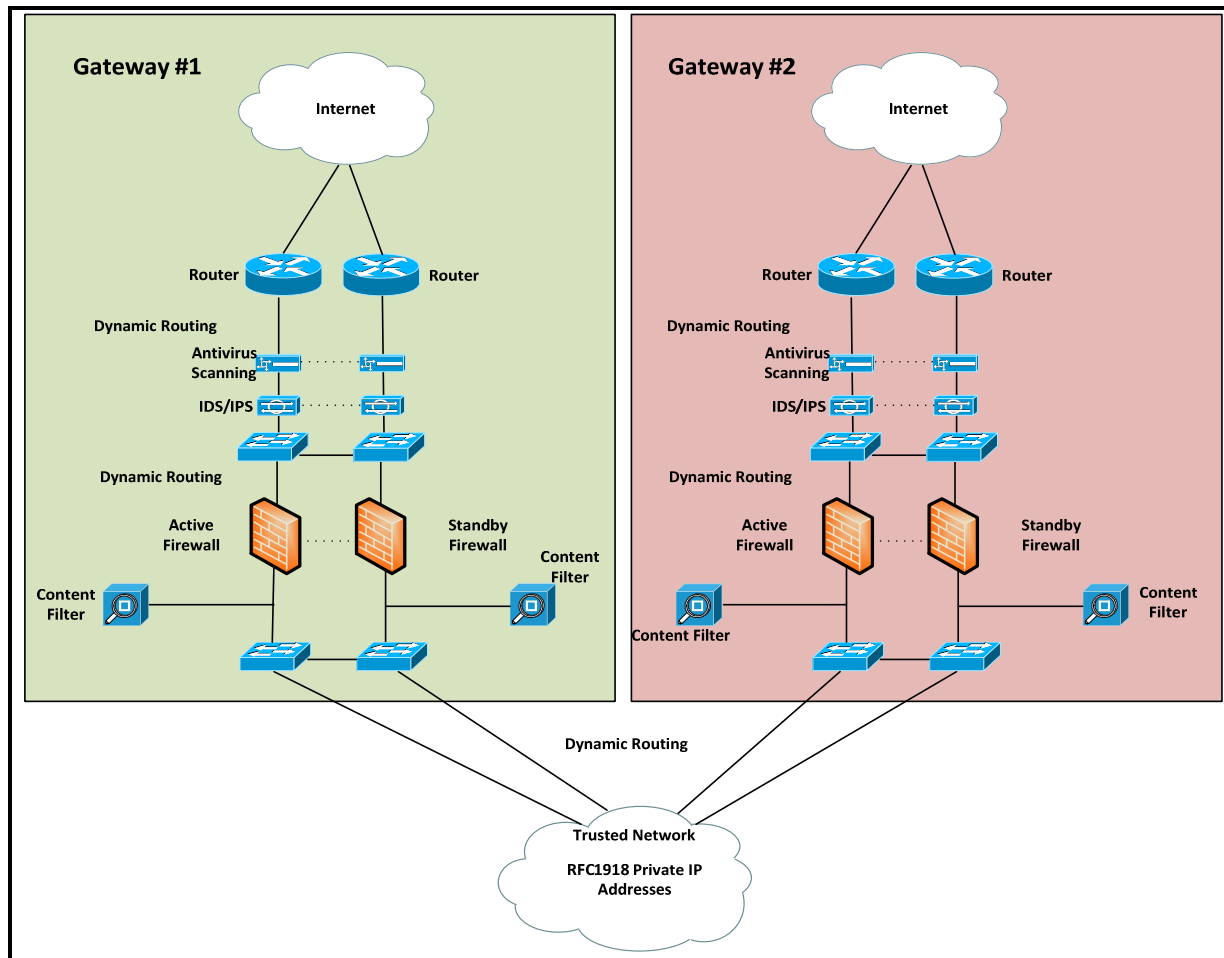


Figure 2.14 Sample Topology: Full Gateway Deployment

A few interesting points about this design. The first is that this is a very common deployment inside of the enterprise today, with multiple devices possibly from different vendors each specializing in their purpose. The cost to implement this from both a capital and expense perspective is costly and the demands for those funds grow as the size and needs grow (Gosal, 2006). The size and type of company will normally dictate many of these needs but also the

company sector could as well. An example of this might be healthcare or a financial company which is bound to protect data in different ways than a typical enterprise. Costs today have crippled organizations from deploying the network perimeter they need and figure 2.14 shows that it may have a direct impact on the scalability of the design. It's possible that a company must forgo using a technology because of a lack of funds to implement. Each security measure must be weighed to see the cost to benefit ratio in addition to the risk factor that should have been identified in the risk analysis.

From the management perspective, most of these devices are managed by either command line interfaces (CLI) or through enterprise management systems (EMS) that are either thick clients that reside on a desktop or through a web services front end. From figure 2.14, if it is assumed that we have the same vendor providing all functions, which would be a best case scenario for management, there still may be several different management techniques that need to be utilized to fully manage this architecture. The router's CLI for example, would be different from the firewall's management application. So best case, there are many different methods for managing this environment. Couple this with the sheer number of management points for enforcement and the architecture begins to show its flaws.

This forays into the operations of the perimeter network which now is quite complex. Contrasting from figure 2.11 to figure 2.14, the amount of support has grown extensively. As was assumed in the previous example that the entire perimeter network is a single vendor, there is still a need to have multiple skilled operations staff members to support the various device functions. Staffing these needs can become costly depending on how different the products are from each other. If the other extreme is taken and it is assumed that each product is a different vendor, the conjecture would be that the staff may need to be skilled in very specialized presence

points. This may be a subject matter expert (SME) for the firewall, a SME for the IDS/IPS and so on. With different touch points, management points, operational staff, the complexity of the environment becomes quite apparent.

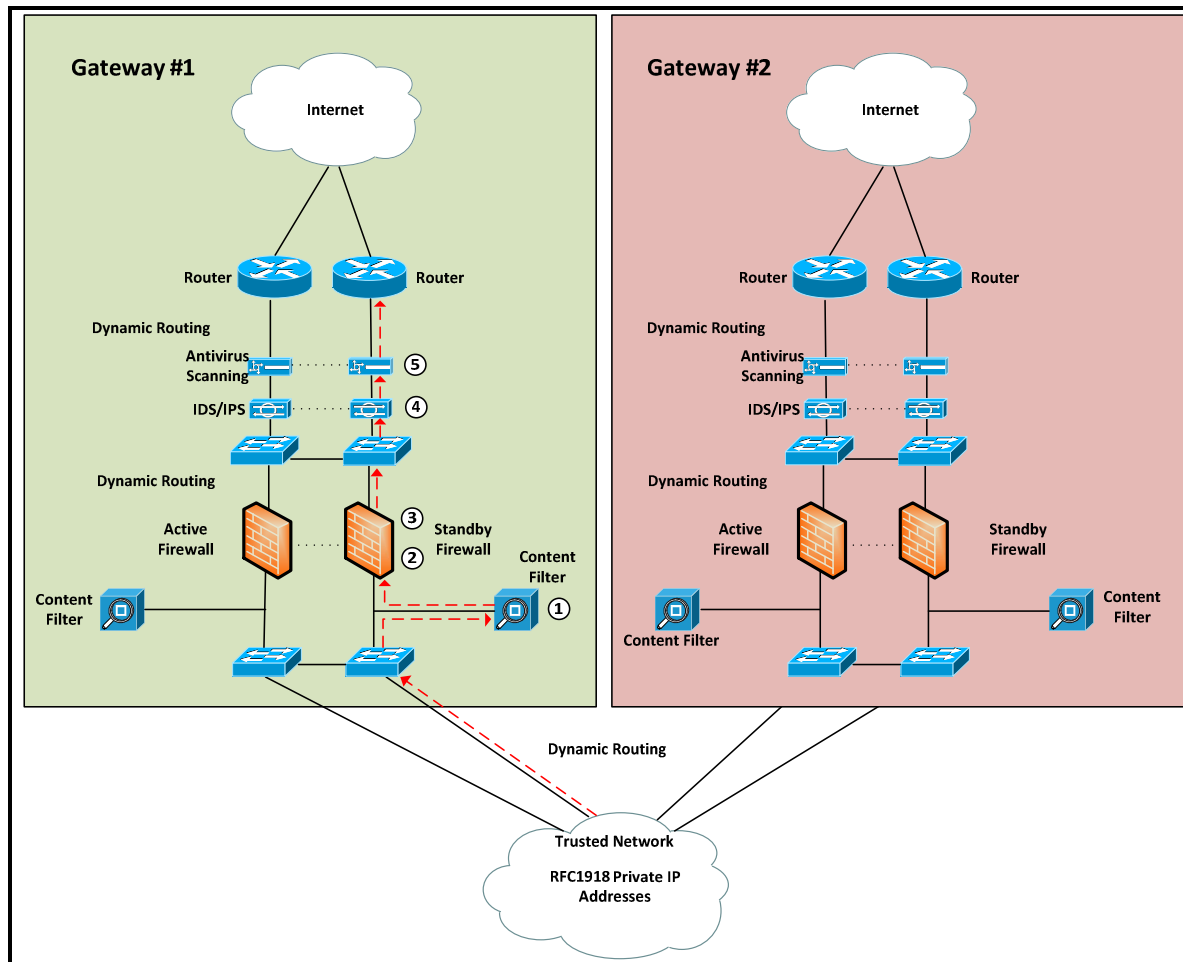


Figure 2.15 Packet Flow Through Full Gateway Deployment

From the perspective of complexity, there is a vital need to look at the flow of traffic through this design. Figure 2.15 attempts to quantify a simple IP packet that must traverse this perimeter network. As routing guides the packet through the maze of threat management products, it is identified that the packet is being observed at each stop. As the packet enters the perimeter, the router will open the packet up to evaluate layer 3 and layer 4 information of the OSI stack. It will then repackage the packet and send it onto the next device in the chain. In this

case, the IDS/IPS and anti-virus devices are inline. These devices will open the packet up as well but are required to dig deeper into the payload of the packet to look for malicious data. This is processor intensive and does introduce latency in the delivery. Once completed, the packet is repackaged and sent to the firewall where it is yet again opened for inspection, only similarly to be forwarded on.

Two interesting notes from this are the traffic latency and inefficiency that is introduced (Fortinet, 2011). Since many of the devices in the path are performing their function in software and not in hardware, there may be significant delay in the overall delivery of each packet. Couple this with the desire for redundancy and resiliency which could come in the form of geo-redundant gateways and the ability to support this environment could challenge the security teams. If there is a problem at any of the points along the way, identifying and finding the problem could be difficult. There is also an inherent reliance upon people from different backgrounds and expertise to work together to keep the system troubleshooting holistic.

A final observation of the existing architectures that are present in the enterprise is one of compliance, reporting and analytics. Security has for some time been focused on how to correlate events. Because when an attack occurs there are normally several flags that if all raised could lead a security professional to quickly understand what's going on. Correlation of the events across different vendors has been an area of the industry that leaders such as Arcsight have attempted to solve. Arcsight is a security information management platform that takes logs from each device and attempts to draw this correlation. In a best case scenario, it may be able to provide a substantial value-add to the perimeter but this comes again with a cost. Without a product such as this, the operations staff is forced to pull logging from disparate locations and

attempt to manually correlate the data. Even on small networks, the amount of data that is captured in the logs of a single device can be overwhelming.

Unified Threat Management

From the previous sections the research has identified issues at multiple layers of the support model that plague the industry's current state of separate and autonomous devices. Figure 2.14 shows just how complex the system can become in an enterprise that has high bandwidth and data processing needs. Scalability of this architecture is possible but at the expense of operational complexity, operational costs, security audits, reporting and overall inefficiency of the traffic.

Unified threat management or UTM as its commonly called is the next evolution of the firewall appliance to utilize application specific integrated circuits or ASICs that are purpose built to offer hardware accelerated speeds for the various forms of risk mitigation. At the core of this UTM model is the next generation firewall or NGFW. The industry realizes that the firewall is the best place to consolidate because it is the core of the filtering of the traffic. The other adjunct techniques are overlapped technology that utilizes this NGFW engine to provide a complete system in form of a consolidated appliance.

UTM is the next evolution of the security perimeter by expanding the focus of the firewall while increasing the firewalls ability to inspect and react to traffic. Because the firewall has been enhanced, it is important to understand why this particular appliance was selected as the core of the UTM model. The NGFW has the ability to leverage breakthroughs in hardware ASICs and network port speeds (Messmer, 2010). In the past where bandwidth needs pushed the firewall to expand into multiple gateways, the hardware available today is able to push higher speeds. With speeds of multi-ten gigabit levels and the addition of ASICs that are adaptable

enough to be modified to new threats, the firewall was the most logical place to start as the core. In looking at traditional firewalls, it's noted that the function of them is to find applications and either permit or deny them. Because of the capabilities of past platforms, they have been restricted to looking at source/destination IP addresses and ports. Well known ports such as HTTP or web traffic are normally configured to use TCP port 80 for example. But these are just generally accepted guidelines and it's true that applications can be run on any port that the application developer desires. So in theory, a web session could be programmed to utilize port 777. With the traditional firewalls, the definition of policies is built upon the idea that applications always use well known ports. Security has proven that this is a major flaw with existing firewalls.

NGFW are now application aware without needing to rely solely on the port numbers. Application aware firewalls are able to look deeper into the packet to find out exactly what application is being utilized, regardless of the port numbering (Messmer, 2010). This awareness allows security perimeter engineers to permit or deny applications like peer to peer clients without needing to painstakingly add a plethora of rules and still not completely mitigate the risk. If there are new applications or new types of traffic that are not recognizable to the firewall, the adaptive nature of the firewall and the soft programming of the ASICs allow the vendor to react quickly with new capabilities.

With the NGFW at the heart of the UTM model, the remaining pieces of the perimeter security are identified and are able to collapse into this high performance platform (Messmer, 2010). Figure 2.16 shows the combination of these elements including routing, IDS/IPS, antivirus and content filtering.

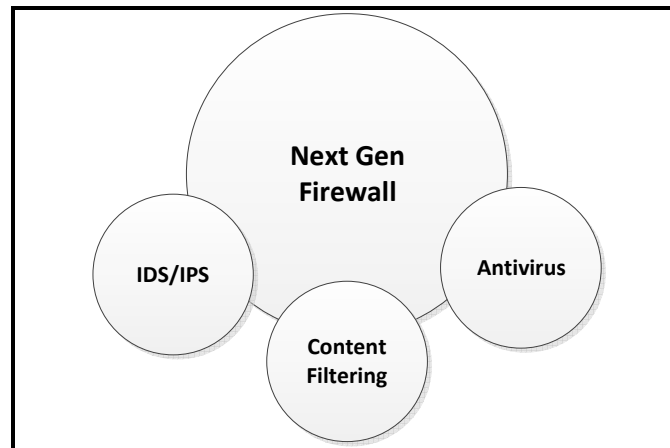


Figure 2.16 UTM Model: Next Generation FW Core

Network functions across the gateway network are driving more routing intelligence into the firewall which is why the collapse of routing functions into the firewall are a perfect example of this convergence. Traditional firewalls have been quite static by nature and relied on network infrastructure around it to steer the packets in the right direction. Today's multiple gateway designs however rely on the firewall knowing the routing topology in order to efficiently route and reroute around failures. In addition to this, an aspect that has been missing in most current firewalls is the idea of quality of service or QoS. As enterprises continue to see surges in IP voice and video, both of which are time sensitive in their delivery, the firewall must be able to accept packets and prioritize them so that the time sensitive protocols are sent out in an expedited fashion. The ability to identify and schedule high priority traffic has always been a weak spot for the traditional firewall. With NGFW, QoS is built into the hardware accelerated data planes that the packets are forwarded on. As the NGFW continues to evolve, hardware vendors are acknowledging that the firewall needs to have similar functionality to the enterprise router which includes the routing, QoS and other technologies such as multicast, which allows for more efficient delivery of packets destined for multiple interested listeners. All of these features, now

being introduced into the NGFW, compliment the edge routers and also provide more advanced traffic shaping which could eliminate the need for additional hardware.

The next security measure that UTM attempts to consolidate into the next generation appliance is the IDS/IPS. IDS/IPS, as outlined in earlier sections, is the function of detecting traffic that matches a pattern of data that is known to be an intrusion attempt or one in progress. Next generation IDS/IPS systems have been evolving just as the firewalls have (Messmer, 2010). Mathematics and statistics have been used in new technology to assist in recognizing and preventing these intrusions immediately and without signatures. The way these platforms work is that they observe the network for a period of time to gather statistics on what the "normal" baseline of the network should look like. This baseline is a believable view to the IDS/IPS of what is safe and normal. If there is a variance in the traffic that is outside of the threshold, the IDS/IPS senses this is an attack and can either notify or proactively shut down the traffic. This new type of technology coupled with traditional signature based detection has brought another efficient mechanism to the UTM model.

Antivirus software relies on similar technology of signature based detection to be effective. Viruses or malware that has been written in the past must be known and uploaded to the antivirus device where it can then detect the malicious data. As with the IDS/IPS appliance, similar strides with computational detection have been getting incorporated into the antivirus devices (Greene, 2007). Leveraging this type of technology becomes advantageous because core functions are similar between the two mitigation techniques. Where IDS/IPS technologies focus on intrusion attempts where the exploiter is trying to gain access to something, antivirus technologies focus on the prewritten code that attempt to propagate and cause service disruption

or destruction of data. While they have different purposes, they protect in very similar ways, so it seems logical that they would leverage similar UTM technology.

The final area of UTM that is gaining traction with convergence into a single appliance is content filtering. Content filtering is the scanning of user traffic to determine whether it is allowed or denied. This traffic is normally web browsing activity. The World Wide Web has several million websites that can be accessed by a user inside of the enterprise perimeter. There are business critical applications that utilize web services in addition to casual browsing sites that are acceptable. There are also sites that carry no business need.

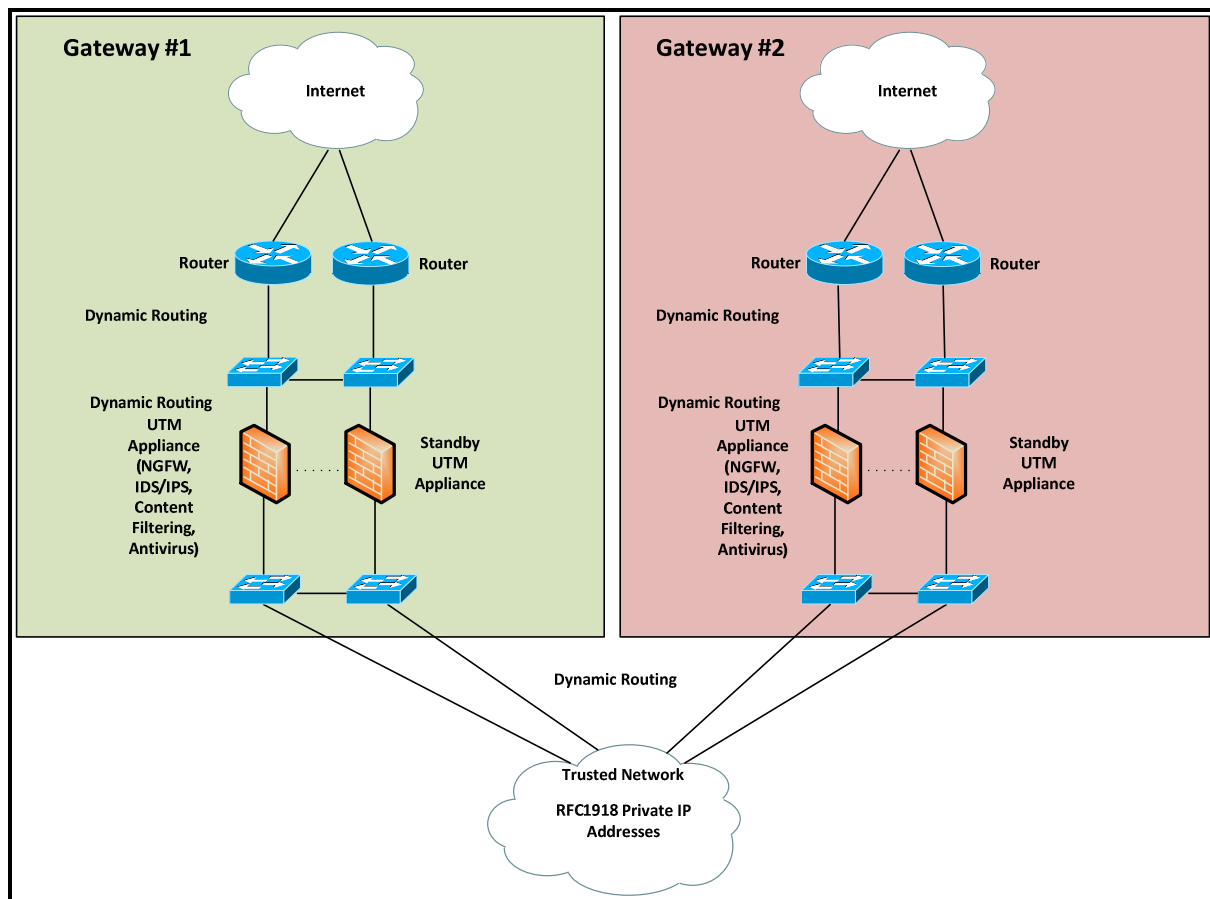


Figure 2.17 UTM Model: Topology of Unified Threat Management

Content filtering parses through the users request for a website and determines based on a configurable basis whether that site is allowed by the company or not. These lists of denied sites are called blacklists. They are updated regularly and reapplied on a continual basis.

Figure 2.17 shows what figure 2.14 would look like with a UTM model applied to it. Figure 2.17 assumes that all functions of UTM are able to be consolidated into a single appliance. The first advantage of this architecture is a single pass, single opening of the packet. In the previous section it was noted that a single packet had to be opened multiple times, once by every device. With UTM and the functions all consolidated into a single hardware based appliance, the packet can be opened one time, have all the security functions perform their analysis and then repackaged and sent on. This in theory should reduce latency and improve traffic efficiency.

Another advantage of the UTM architecture is obviously the reduction in the amount of hardware involved which will drive the overall cost of implementation and support down. Consolidated management of the various functions can result in a cleaner, single pane of glass view into the perimeter that allows less touch points for management. Things that were not easily accomplished now have more promise in this architecture such as correlation of events. Since a single appliance is inspecting and observing all different security postures, the vendor can more easily correlate the triggers between them. Reporting and logging now have a more consistent and uniform appearance. The overall design from a support perspective and complexity of the packet flow are substantially reduced. These advantages along with new abilities such as centralized identity management are making their way to UTM devices. Identity management allows the company to track traffic sources to specific individuals with the use of applications like Microsoft Active Directory.

UTM is currently being incorporated into major security vendors' equipment and tweaked to provide market leaders in this space. Gartner has a magic quadrant for these next generation security appliances. The thesis focuses on evaluating the reality of two such UTM appliances compared and contrasted with a traditional firewall appliance. The results provide interested parties with key points that differentiate the vendors from each other and how the evolving market for UTM is attempting to meet the ongoing challenges with the perimeter security architecture.

Chapter 3 – Methodology

Introduction to Methodology

The study will use mixed methodology. Its purpose is to reach into two different points of research in order to provide color to the "why" and "how" of the UTM approach. The first objective is to investigate why the security industry is in this position with respect to network perimeter security. This investigation elaborates on many of the areas discussed above but in more detail to provide a clear backdrop for how the enterprise security environment currently has this architecture. It is vital to understand what the current state is so that the benefits of UTM can be quantified.

The second research objective is to investigate what the vendors are doing about the current state of security posture by implementing and developing UTM features. With this research objective, several vendors will be analyzed, compared and contrasted against each other in order to show clearly where UTM is effective and where there are still possible shortcomings. This research point is important to show how the market, specifically the vendor, is responding to the current challenges with today's security perimeter and how next generation technologies within their product suites will provide UTM functionality to meet tomorrow's demands.

The research area has been narrowed down to focus on unified threat management as one area of the tiered security architecture. Knowing there are several layers to the security posture of an organization, the focus on this area provides an in-depth look at the perimeter security where organizations place most of their emphasis (Northcutt, Zeltser, Winters, Frederick, & Ritchey, 2003). It is further refined to include only a subset of what encompasses unified threat management. As indicated above, unified threat management consists of many different types of security protections. For the purposes of staying grounded, the research only investigates four of

the core technologies, those being firewall, intrusion prevention, content filtering and malware mitigation.

Method

From the research objectives, a single methodology did not work to meet the goals of the thesis. In the thesis a mixed research methodology was used that utilized both investigative and design science. It was important to understand the shortcomings of existing security architectures and uncover how these have shaped the evolving UTM technologies. The investigative portion of the research is augmented by illustrating through vendor comparisons, the strengths and weaknesses with products that exist in the market today.

Design Science Research

Design science is concerned with the analysis of a problem and potential solutions that may exist to produce an artifact to solving this problem. Because the cycle of design science is iterative, we see that the process to complete this has some basic starting points but is primarily concerned with continually refining the solution to produce better or more efficient ones. In the case of UTM, design science is appropriate because the problems with current network perimeter security are quite evident. Problem definition being the first step, the research has analyzed what the current state of enterprise network security perimeters is and how it is flawed. This deep research into the various components and how they operate today drives a hypothesis that indicates that other solutions are better able to meet this demand. For the purposes of this thesis, UTM will be our focal point in proving the hypothesis that a combined architecture of security solutions will solve our problems.

Evaluation

Prototyping and modeling are two very effective methods of analysis of the hypothesis. In our case, the UTM field is still evolving and full scale prototyping of every feature that UTM offers is not an easy possibility. Because of this, the research will be a combination of author generated lab evaluation coupled with industry research. Modeling will consist of a clear set of requirements of which each vendor will be evaluated to see at what degree they can satisfy the potential solution, which again is primarily focused on the four technologies of UTM. Constructs will be used to propose ideas based on the research that indicates that a certain construct will solve the intended problem. The collection of these constructs will become the foundation for the model that is created.

The primary output from this part of the research will be the actual evaluation of the technologies by various vendors. The evaluation will be critical in drawing a conclusion about the UTM solutions. A matrix is used to subdivide the UTM areas of focus into various categories for cross vendor comparison. The research shows the advantages and disadvantages of each category as it relates to vendor capabilities.

Figure 3.1 is the topology used in lab testing for the results section. It consists of a 12mbps DSL link to the Internet and publicly assigned address space to provide global access to the lab. The outside multilayer switches are used primarily to create the virtual LANs needed to test the vendor equipment in addition to interaction with the routing of the environment.

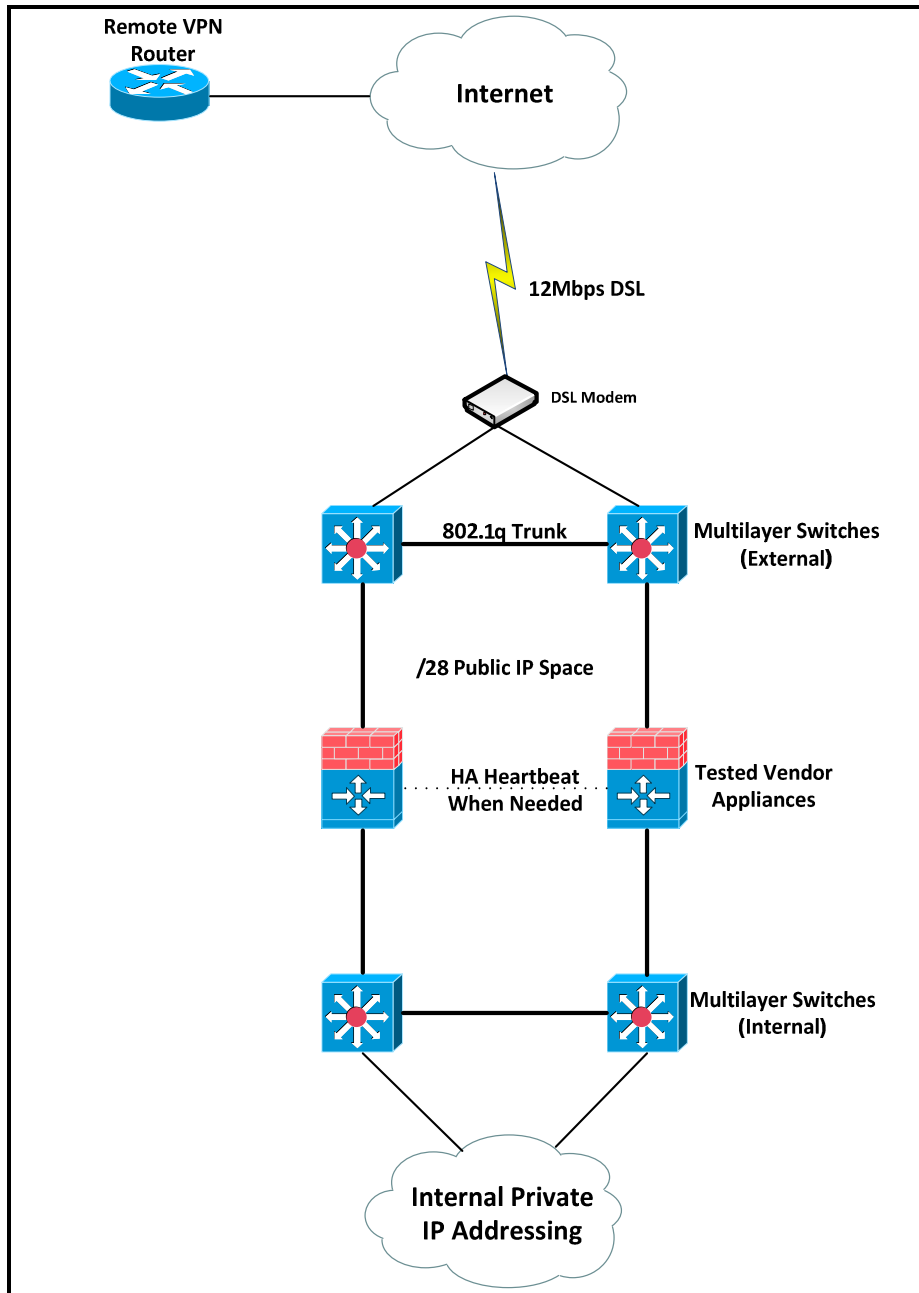


Figure 3.1 Lab Topology For Testing

The vendor test equipment was placed into the environment with the highest configured bandwidth which in all testing was 10Gbps. The internal multilayer switches were used in the same manner that the external switches were with the exception of any testing of traffic from a

user perspective was accomplished here. This allowed for testing of basic firewalling, content filtering, antivirus, NAT and other elements of the test.

Chapter 4 –Results

Introduction to Results

Firewalls have become the foundation of enterprise security. Historically they have been funneling points of traffic to be scanned for potential risk. As the communication of traffic between individual protected networks continues to grow, the demand on these devices for both bandwidth and functionality has grown. The results from the research conducted are represented in this section by a comparison matrix that attempts to not only differentiate the traditional firewall platform from the unified threat platforms but also show some comparison between competing products in this space. The matrix will provide a snapshot of where the traditional firewall has evolved to meet the limitations described in earlier sections. It also shows where these next generation platforms are meeting the demands and where they might still be falling short.

The traditional firewall has been around for nearly 25 years and many enterprises are seeing these new products being released during a time of refresh or cyclical reengineering of their perimeter. For this reason, it's also valuable to take the research garnered from the matrix and apply it to decision making processes today with respect to life cycle of equipment. The matrix should provide an idea where these products excel and where they fall short.

Matrix Results

The evaluation matrix attempts to quantify and explain the similarities and differences between the traditional firewall and the UTM platform while also drawing out differences between two market leaders in this space. The results will consist of lab testing by the author coupled with backing information from industry research through Gartner and NSS Labs. Because the UTM market is relatively new, Gartner does not have enough research compiled to

conduct a readout but the core of the UTM market being the NGFW does. Figure 4.1 shows the Gartner group's ranking of the NGFW quadrant. The quadrant is broken down into four main areas but for our research, we will focus on the "leader" and the "visionary" sections. Leaders are well established dominant vendors in this technology area while visionaries are companies who are innovating in ways that the leaders are not. The interesting comparison here is that while much can be learned from the market leaders, there is a lot of value in evaluating the visionaries as they have been noted by Gartner to be pioneering new features and/or functionalities.

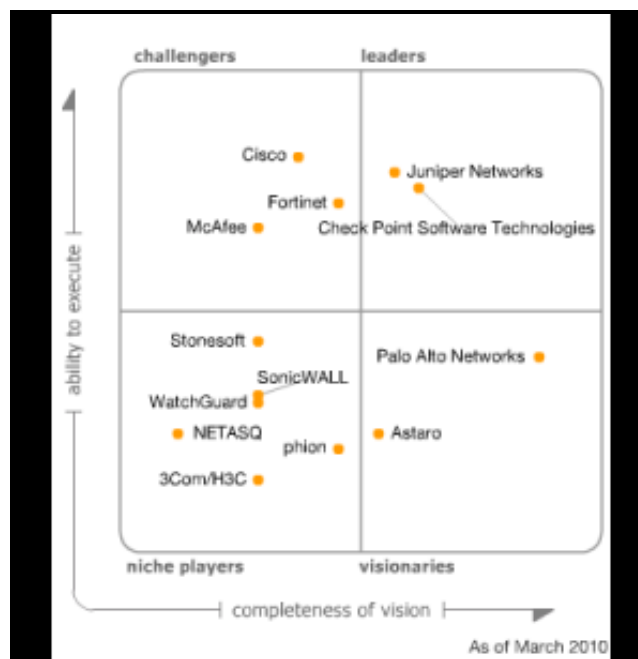


Figure 4.1 Magic Quadrant For Enterprise Network Firewalls (Young & Pescatore, 2010)

For the results of the matrix, Check Point has been chosen as the vendor in the leader quadrant, specifically the Power-1 11067 chassis. For the challenger quadrant the Palo Alto PA-5060 will be evaluated. Finally for the comparison with a traditional firewall, we will use the Cisco Adaptive Security Appliance 5585-X. Table 4.2 is the matrix showing each respective

Table 4.2 Unified Threat Management Evaluation Matrix

Unified Threat Management Evaluation Matrix							
	Evaluation Weight	Check Point Power-1		Palo Alto PA-5060		Cisco ASA 5585	
		Score	Weighted Score	Score	Weighted Score	Score	Weighted Score
Functionality							
Routing	3	3	9	4	12	3	9
Packet inspection	5	4	20	5	25	4	20
NAT	3	4	12	3	9	5	15
VPN	3	4	12	3	9	5	15
Voice/Video Support	4	4	16	2	8	5	20
Content Filtering	5	1	5	5	25	0	0
Antivirus	5	1	5	4	20	0	0
Application Identification	4	2	8	5	20	1	4
IPS/IDS	5	4	20	3	15	0	0
Virtualization	4	2	8	4	16	1	4
High Availability	3	4	12	3	9	3	9
Quality of Service	4	2	8	4	16	2	8
Operations							
Unified Management	5	5	25	4	20	2	10
Unified Logging	4	5	20	5	20	2	8
Command Line Interface	1	3	3	4	4	4	4
Policy Conversion	1	3	3	4	4	4	4
Misc							
Education	2	4	8	3	6	3	6
Support	4	2	8	4	16	3	12
Cost	5	2	10	4	20	3	15
Totals							
<i>Totals of Points/Weights</i>	<i>70</i>	<i>59</i>	<i>212</i>	<i>73</i>	<i>274</i>	<i>50</i>	<i>163</i>
<i>Weighted Percentages</i>		<i>61%</i>		<i>78%</i>		<i>47%</i>	

<i>Total Points Possible</i>	95				
<i>Total Weighted Points Possible</i>	350				

category that was evaluated with a category weight, vendor score and vendor score with weight applied. At the bottom of the matrix are the totals of each vendor's score and their weighted percentage. Table 4.2 was a result of hands on lab testing of each platform coupled with industry research from Gartner and NSS Labs.

Evaluation of the Matrix

In order to understand how the numbers in table 4.2 were achieved, there must be some clarity around the category, why the weight was added and how each vendor scored with relation to that category. The evaluation matrix attempts to quantify and explain the similarities and differences between the traditional firewall and the UTM platform while also drawing out differences with each vendor compared. Each area of this snapshot will be discussed in detail in the subsequent sections.

Functionality

This section primarily focuses on how the platforms actually provide technical features to meet requirements in the areas outlined in this thesis. These areas, captured in table 4.2, are a combination of what is expected from traditional firewalls in addition to the features that the UTM appliances are driving into the market.

Routing



Routing is the process of receiving packets into one interface, looking inside of the layer 3 portion of the header and making a decision about what interface will send the packet closer to the destination. The evaluation of routing has a few different criteria. The first is the mode that

the firewall can operate in. The types of modes are layer 2, layer 3 and virtual wire. With a firewall operating in layer 2 mode, the firewall acts as a layer 2 Ethernet switch where it will be part of the MAC forwarding plane. In this mode, the Ethernet frames are forwarded up to the firewall where policies and UTM functions can be performed. Because the firewall is acting as a switch, there is no noticeable “hop” in the flow of packets.

In layer 3 mode, the UTM device acts as a true router carrying a full routing table and using dynamic protocols to discover routes to destinations. These dynamic protocols allow the UTM device to interoperate with traditional routers to be a part of the topology. This helps with redundancy and resiliency. Protocols such as RIPv2, OSPF, BGP and static routing were all evaluated.

Table 4.3 Functionality – Routing

Functionality - Routing							
	Check Point Power-1	Palo Alto PA-5060	Cisco ASA 5585	Validated By			
Deployment Modes	L2, L3	L2, L3, Virtual Wire	L2, L3				
Routing Protocols Supported	RIPv2, OSPF, BGP, Static	RIPv2, OSPF, BGP, Static	RIPv2, OSPF, EIGRP, Static				
Policy Based Forwarding	Not Supported	Supported	Not Supported				
VLAN Support	1,024 VLANs	4,094 VLANs	1,024 VLANs				
Aggregate Links	Supported	Supported	Supported				
Multicast Support	IGMPv2/v3, PIM-SM/DM	Not Supported	IGMPv2/v3, PIM-SM/DM				
IPv6 Support	Supported	Supported	Supported				
Score	3	4	3				

 Evaluated in Researcher's Lab
 Evaluated by NSS/Gartner

Virtual wire mode is a physical layer technology also called “bump in the wire” where the UTM device is not visible by any means other than being placed between two endpoints who

believe they are directly connected to each other. This is particularly valuable when you do not need to perform any NAT, create any VPN connections or any other feature that requires that the UTM device terminate a connection.

Other aspects of routing that were evaluated were the support of not only unicast but multicast as well. Multicast is the idea of sending packets to a group IP address and interested parties subscribe to listen to the stream. It has advantages over unicast in that if there are multiple parties interested in the same information, the packets are not duplicated across the network. In the past UTM devices and firewalls have not supported multicast.

The last feature evaluated is the support of the next generation of IP with version 6. The current version of IP is IPv4 which is showing serious signs of exhaustion for globally unique addresses. IPv6 is the next iteration of IP which allows for unprecedented scale so supporting this is an absolute must out of any UTM device. Table 4.3 shows that all of the platforms evaluated support layer 2 and layer 3 modes however it should be noted that Palo Alto supports an additional mode that the other two do not which is referred to as virtual wire. Virtual wire allows the UTM appliance to operate at layer 1 of the OSI model, which would be simply passing packets. The devices on both sides of the UTM appliance believe they are connected directly to each other however; the firewall intercepts traffic in the flow for inspection. In this mode, the firewall cannot perform certain functions such as NAT or VPN termination.



Table 4.3 also shows that all of the appliances are compatible with most of the industry standard protocols for routing. Check Point and Cisco take the lead in the fact that it can support multicast traffic which Palo Alto cannot. Overall with the scalability of VLANs, support of policy-based routing and the additional deployment mode, Palo Alto scored higher in the tests for routing.

Packet Inspection

Since the core of the UTM model is the next generation firewall, the importance of the UTM device being able to not only perform what traditional firewalls have been doing for years but also improve upon stateful inspection and speed of passing packets is critical. In the details of this category, the vendors were evaluated on raw speed at which they can parse their rules to permit or deny a packet. Because all three vendors have been perfecting their firewall engines to handle common things like spoofing, session hijacking and other IP based attacks, the primary criteria being evaluated here is pure performance of the firewall engine in throughput while in protect mode. Table 4.4 really starts to show how the two UTM appliance begin to differentiate themselves from the traditional firewall in that of the Cisco ASA.

Table 4.4 Functionality – Packet Inspection

Functionality - Packet Inspection							
	Check Point Power-1	Palo Alto PA-5060	Cisco ASA 5585	Validated By			
Firewall Throughput	20Gbps	20Gbps	10Gbps				
Maximum Connections	1,200,000	4,000,000	4,000,000				
Connections Per Second	58,000	120,000	200,000				
DDoS Support	Supported	Supported	Supported				
SSL/SSH Decryption	Supported	Supported	Not Supported				
Authentication	Supported	Supported	Only Supported for VPN				
Single Pass Inspection	Supported	Supported	Not Supported				
Score	4	5	4				

	Evaluated in Researcher's Lab
	Evaluated by NSS/Gartner

Raw firewall throughput was evaluated to show that both Check Point and Palo Alto have advanced their ASICs to push the inspection limitations up to the 20Gbps realm. The Cisco ASA, at half of that rate, also shows its age with the failure to meet features such as SSH/SSL

decryption on the fly to inspect packets in addition to a lack of supportable identification of user traffic. Because of these things, the ASA is able to edge out the UTM devices on connections per second and total numbers, but only at the expense of the lacking features. This is the first hint in the UTM testing though that enabling all of the “bells and whistles” will come at a tradeoff with overall performance.

The most important element to note from this testing is that with these advanced features, the two UTM appliances are able to support a single pass inspection. This means that while the ASA and other devices would have to open the packet several times to evaluate a similar features, the UTM appliances are able to remain efficient by opening the packet one time for the application of rules. With the overall features supported, high connection limit and the high throughput, Palo Alto scored the highest in this area.

NAT

NAT refers to the process of taking one IP address and changing it in the IP header to another. This may be needed in order to connect a privately addressed network to another network such as the Internet. It can also be used to provide access from another network into your private network. Translations originally were a way to save globally unique IPv4 addresses but have been used over the years an added security benefit as it hides the topology of the private network. The evaluation of the ability to NAT a packet comes in terms of how many translations a UTM device can support and how quickly it can process these types of requests.

Table 4.5 shows once again where the traditional firewall excels at what it has been known for. The Palo Alto is limited to a finite number of translations while the Check Point and Cisco platforms are only bound by the limitations on the memory that the NAT table is held in.

Table 4.5 Functionality – NAT

Functionality - NAT

	Check Point Power-1	Palo Alto PA-5060	Cisco ASA 5585	Validated By
Maximum NAT Sessions	Bound by Memory	250,000	Bound by Memory	
NAT Modes	1:1, N:N, M:N	1:1, N:N, M:N	1:1, N:N, M:N	
NAT Types	Dynamic, Static	Dynamic, Static	Dynamic, Static	
Enhanced NAT Functions	Limited Support	Limited Support	Supported	
Score	4	3	5	

	Evaluated in Researcher's Lab
	Evaluated by NSS/Gartner

All three devices support the same modes and types of translations but the Cisco ASA has the ability to support some enhanced NAT functions such as subnet to subnet translation and application layer translations with relative ease. Where these features are simply one or two commands in the Cisco ASA, they are either not supported at all or are cumbersome to configure on the UTM devices. In the case of network address translations, the traditional firewall came out on top with scoring.



VPN

Virtual Private Networks or VPNs logically extend the borders of the enterprise network by using encryption and routing over networks that are not necessarily controlled by the enterprise. VPNs are established between endpoints that form a logical tunnel with each other and through a systematic process of credential exchanges, form a secure connection between the two. VPNs are a popular way to extend the enterprise network in a secure fashion and are known for a quick and cost efficient alternative to provisioning physical leased line circuits. For the purposes of the evaluation, the vendors were assessed on how many VPN connections they

can support, at what throughput and also what types of VPN technologies were completely interoperable.

Table 4.6 Functionality – VPN

Functionality - VPN							
	Check Point Power-1	Palo Alto PA-5060	Cisco ASA 5585	Validated By			
IPSEC VPN Throughput	3.7Gbps	4Gbps	4Gbps				
IPSEC VPN Max Tunnels	Bound by Memory	8,000	10,000				
Routing Over VPN	Only supported with additions to the Operating System	Supported	Supported				
VPN Compatibility	High	Medium	High				
Score	4	3	5				

	Evaluated in Researcher's Lab
	Evaluated by NSS/Gartner

As illustrated in table 4.6, Palo Alto is quite cautious about the total number of VPNs that it supports while the Check Point UTM device allows for as many as the memory can hold. Obviously tweaking the memory will produce varied results. Utilizing dynamic routing protocols over the VPN were recognized in both the Palo Alto and the Cisco ASA but were only available in the Check Point appliance with some operating system work and were not readily available in the UTM application. Once again, a tried and true feature like VPN has been perfected by the one of the market leaders from the traditional firewall realm in the Cisco ASA which scored higher points in this area.

Voice/Video Support



Enterprises are quickly moving to IP based voice and video solutions such as Voice over IP and video conferencing. By nature, extending the reach of these technologies outside of the

enterprise network is growing so the UTM device must be capable of supporting this requirement. Voice and video can be subdivided into two different functions; signaling and media stream. The first function of signaling involves how the endpoints find, negotiate and setup a call with another endpoint. Common protocols in this space are SIP, H.323 and SCCP. The second part of the equation after the call is setup is the actual media stream that would represent the voice and/or video. This is traditionally RTP or SRTP packets. In both cases, the UTM device must be able to understand and pass the signaling and the media stream through its protection mechanisms. Because quality of service is evaluated in a later section, the main point here is the devices ability to handle many different interpretations of the various signaling protocols. This is depicted in table 4.7.

Table 4.7 Functionality – Voice and Video

Functionality - Voice & Video Support

	Check Point Power-1	Palo Alto PA-5060	Cisco ASA 5585	Validated By
SIP Support	Supported - With Special Release	Supported - RFC Only	Supported	
MGCP Support	Supported	Not Supported	Supported	
H.323 Support	Supported	Supported	Supported	
SCCP Support	Supported	Supported	Supported	
NAT'd SIP Support	Not Supported	Not Supported	Supported	
Score	4	2	5	

 Evaluated in Researcher's Lab
 Evaluated by NSS/Gartner



The score that Cisco attained in this area is not shocking as Cisco Systems has a lot of history with helping to define these standards but it was a little surprising that Palo Alto had difficulty with an industry standard such as SIP in a NAT scenario.

Content Filtering

Content filtering has been discussed at great length in the previous sections but the general idea is to have the ability to filter primarily web based traffic. When users browse the Internet there is a strong desire to be able to enforce rules about where they can browse to and where they cannot. Another feature that enterprise security teams are looking for is the ability to identify a user by IP address and more importantly by some type of login credential such as active directory or LDAP. In this space, table 4.8 shows the evaluated effectiveness of the categories that each vendor allows for blacklisting sites, ease of configuration and the ability to identify users.

Table 4.8 Functionality – Content Filtering

Functionality - Content Filtering								
		Check Point Power-1		Palo Alto PA-5060		Cisco ASA 5585		Validated By
Category Based URL Lists		Supported - Optional Addon		Supported		Not Supported		
Customized Categories		Supported - Optional Addon		Supported		Not Supported		
Customized Block Pages		Supported - Optional Addon		Supported		Not Supported		
Dynamic URL Filtering		Supported - Optional Addon		Supported		Not Supported		
Identity Mangement		Supported - Optional Addon		Supported		Not Supported		
Score		1		5		0		

 Evaluated in Researcher's Lab
 Evaluated by NSS/Gartner

The first one of the real core elements of UTM did uncover some surprises in testing as shown in table 4.8. Content filtering is supported in the Check Point device but only via an optional add-on blade which requires a hardware card and associated licensing. Because this was not part of the base package of the device, the device was scored lower. The Cisco ASA does not support any content filtering as Cisco relies on their IronPort standalone product to meet

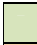

this requirement. The Palo Alto shines with this feature providing easy to configure and an operationally friendly interface. Updated subscriptions from a third party or custom written URL blocks are allowed as are customized responses back to users. The ease with which it is to set up the device for content filtering in the same management plane as the firewall rules made the configuration straight forward. For these features, Palo Alto was awarded the full five points in this area.

Antivirus

Blocking malware is a critical part of the protection mechanism of the security perimeter. The ability for a UTM device to have predefined parameters that are able to catch these malicious programs before they enter the enterprise certainly adds to the unifying theme of the platforms. Evaluating the types of antivirus protection and the speed at which the platform can perform this function is important in determining if the UTM device is an appropriate place to perform this scanning or if stand alone devices are still a better choice.

Table 4.9 Functionality – Antivirus

Functionality - Antivirus							
	Check Point Power-1		Palo Alto PA-5060		Cisco ASA 5585		Validated By
Threat Prevention Throughput	10Gbps - Optional Addon		10Gbps		Not Supported		
Application threat prevention	Supported - Optional Addon		Supported		Not Supported		
OS Threat Prevention	Supported - Optional Addon		Supported		Not Supported		
Stream based scanning	Supported - Optional Addon		Supported		Not Supported		
Sypware	Supported - Optional Addon		Supported		Not Supported		
Viruses	Supported - Optional Addon		Supported		Not Supported		
Worms	Supported - Optional Addon		Supported		Not Supported		
Score	1		4		0		

 Evaluated in Researcher's Lab
 Evaluated by NSS/Gartner

Once again the traditional firewall does not support antivirus scanning and that Check Point offers this feature but only at the expense of an add-on. For this reason and again the ease of configuration of the Palo Alto, their score reflects dominance in this category.



Application ID

The identification of applications has become important as more and more applications can be run on any TCP or UDP port. This particularly has been apparent in software such as peer to peer software that will hide itself behind well known ports such as port 80 which belongs to web traffic. It is not acceptable anymore to simply scan for ports. The devices must dig into the layer 7 part of the packet to determine what application is actually being evaluated.

Application identification is a core strength of evolving NGFW and table 4.10 shows the strength of each vendor in this space.

Table 4.10 Functionality – Application Identification

Functionality - Application Identification						
	Check Point Power-1		Palo Alto PA-5060		Cisco ASA 5585	Validated By
Identification of Applications	Supported - In Software		Supported - Hardware		Limited Support - NBAR	
Application ID in SSL	Supported - In Software		Supported - Hardware		Not Supported	
Application ID in SSH	Not Supported		Supported - Hardware		Not Supported	
Application Based Traffic Shape	Not Supported		Supported - Hardware		Not Supported	
Score	2		5		1	

 Evaluated in Researcher's Lab
 Evaluated by NSS/Gartner

As applications continue to break the rules of de facto standards for port assignments, the identification of applications by means of the actual payload is becoming more a necessity.

Couple this with the emergence of applications that are not business critical that will probe for any open port to use to get themselves outside of the perimeter and you can see how important application identification is. Simple identification of the application is supported by all three platforms but each to a varied extent. In Check Point, they support a high degree of application identification but it is done in the Check Point software whereas the Palo Alto has soft reprogrammable ASICs that are used to find this traffic at near wire speeds. The Cisco ASA supports a rudimentary form of application identification with network based application recognition or NBAR for some time but it is also done in software and is limited to an isolated set of protocols.

Palo Alto has numerous mentions in the industry for this feature which allows administrators the ability to instruct the UTM appliance to block peer to peer file sharing, regardless of what port it is running on. The dynamic nature of this search and destroy mentality clearly points out Palo Alto is a leader in this space.



IDS/IPS

As discussed above, the IDS/IPS feature is designed to thwart attacks that attempt to gain access to key devices inside of the enterprise. Differentiation in this space is how the IDS/IPS system works by either subscription or mathematical algorithms and how much the process of turning on IDS/IPS features affects the raw performance of the platform. Table 4.11 shows the nature of each vendor with respect to intrusion protection. Check Point is the clear leader in this category as it relates to UTM. The Cisco ASA does not support this feature as Cisco relies on a standalone platform to compete in the category. Palo Alto, although keeping up with Check Point, did not provide as much granularity with respect to configurations. They also did not

support behavior based detection which allows the UTM appliance to learn from previous traffic patterns.

Table 4.11 Functionality – IDS/IPS

Functionality - IDS/IPS							
	Check Point Power-1	Palo Alto PA-5060	Cisco ASA 5585	Validated By			
IDS/IPS Throughput	10Gbps	10Gbps	Not Supported				
Signature Based	Supported	Supported	Not Supported				
Anomaly Detection Based	Supported	Supported	Not Supported				
Behavior Based	Supported	Not Supported	Not Supported				
DOS Mitigation	Supported	Supported	Not Supported				
Customized Signatures	Supported	Supported	Not Supported				
Score	4	3	0				

 Evaluated in Researcher's Lab
 Evaluated by NSS/Gartner

Virtualization


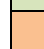
Virtualization inside of the enterprise is not a new concept as companies such as VMware and Microsoft have been performing this function for many years. Network devices such as multilayer switches have been demonstrating virtualization in the LAN through VLANs for several years. In the context of the UTM device, it has become advantageous to virtualize the security appliance. This term means different things to various vendors. Some vendors see virtualization as simply allowing logically separated rule sets. Others believe virtualization is only true in the idea that several completely separated firewalls can be created virtually out of one physical device.

The idea of virtualizing a UTM appliance comes down to the right fit for the right situation. If there are multiple needs for the UTM device and a requirement to separate the

device into different logical threat management appliances, Palo Alto had the most flexibility. It also came with the most base level licenses out of any of the vendors evaluated. Once again, the traditional firewall does not support this type of feature. Table 4.12 outlines the findings from this category.

Table 4.12 Functionality – Virtualization

Functionality – Virtualization							
	Check Point Power-1	Palo Alto PA-5060	Cisco ASA 5585	Validated By			
Security Zones/Contexts	Supported	Supported	Supported				
Virtual Routers	Supported - Separate Platform	Supported	Not Supported				
Virtual Systems	Supported - Separate Platform	Supported	Not Supported				
Score	2	4	1				

	Evaluated in Researcher's Lab
	Evaluated by NSS/Gartner



High Availability

Availability is one of the key components of any network. Availability inside of the security perimeter is paramount. If a UTM device fails, it's critical that the technology allows for seamless failover. High availability refers to the act of having an alternative device available to take over in the event that the primary device fails. The challenge in this area is that because the UTM device is maintaining state awareness for each flow that it is servicing, failover to another device could be disruptive if that device does not have the same state information. In that case, the traffic would failover but any session that is connection based could be disconnected forced to reestablish.

Check Point has been an early pioneer in the area of availability. Table 4.13 shows that they support all of the common availability modes in addition to the clustering of firewalls to allow load balancing. This important feature places the Check Point UTM platform as an edge winner in this category.

Table 4.13 Functionality – High Availability

Functionality - High Availability							
	Check Point Power-1	Palo Alto PA-5060	Cisco ASA 5585	Validated By			
High Availability Supported	Supported	Supported	Supported				
Active/Standby Supported	Supported	Supported	Supported				
Active/Active Supported	Supported	Supported	Supported				
Load Balancing/Clustering	Supported	Not Supported	Not Supported				
Score	4	3	3				

-  Evaluated in Researcher's Lab
-  Evaluated by NSS/Gartner

Quality of Service

Quality of Service or QoS is absolutely necessary based on time sensitive applications such as voice and video. As more of this type of traffic passes through the perimeter, the need to schedule and give priority of forwarding to these applications increases. In order to prioritize traffic so that it is expedited through the chassis, QoS is at minimum a must but also has to be granular enough to control, so that protections can be put in place to limit bandwidth of certain types of traffic as well.



Table 4.14 shows the evaluation of the QoS features of each platform and there is no surprise that each device supports basic network layer QoS. This is enough to expedite the forwarding of traffic that is correctly marked in the Type of Service or TOS bits of the IP header.

Palo Alto shines in this category as, once again it can dig further into the packet and seek out applications without markings. For example, if a voice over IP phone failed to mark the TOS bits correctly, the Palo Alto UTM appliance could still be instructed to find voice traffic and give it priority. It can also place priority from one virtual system to another. This is not possible with the other vendors because of their relative lack of virtualization to this level.

Table 4.14 Functionality – Quality of Service

Functionality - Quality of Service

	Check Point Power-1	Palo Alto PA-5060	Cisco ASA 5585	Validated By	
Layer 3 QoS	Supported	Supported	Supported	Evaluated in Researcher's Lab	Evaluated by NSS/Gartner
Layer 7 QoS	Not Supported	Supported	Not Supported	Evaluated in Researcher's Lab	Evaluated by NSS/Gartner
Low Latency Queues	Supported	Supported	Supported	Evaluated in Researcher's Lab	Evaluated by NSS/Gartner
Virtual System QoS	Not Supported	Supported	Not Supported	Evaluated in Researcher's Lab	Evaluated by NSS/Gartner
Score	2	4	2		

-  Evaluated in Researcher's Lab
-  Evaluated by NSS/Gartner

Operations

One of the key components of the success of UTM relies on its ability to simplify the operations of the perimeter network. With multivendor devices operating in disparate manners currently, operational organizations have a difficult task in not only being proficient on many different management systems but also have challenges with tracking packet flows. Operational criteria for the evaluation of these UTM devices constitute looking at a few key areas. The first would be the actual management interface into each device. Secondly it is important to evaluate how well each vendor has unified the logging and correlation of the various components of the UTM system. Because some systems have special parameters that can only be changed via the

command line interface or CLI, some time was spent in walking through the ease of each vendors CLI to use. Lastly, an important but little talked about feature of UTM devices comes with policy conversion. Chances are if an enterprise is migrating toward a new UTM model, the vendor may not be the same. If the vendor is the same, it's possible that the new UTM policies are substantially different than the traditional format. Because of this, evaluation of each vendor's ability to convert policies from other platforms was taken into consideration. The following sections outline the results of these operational areas.

Unified Management

The management of the UTM device is one of the most important operational aspects of the system. If the security operators cannot easily add, change or delete something quickly and intuitively, the time to react to a risk could start to climb and affect the overall efficiency of the system.

Table 4.15 Operations – Unified Management

Operations - Unified Management

	Check Point Power-1	Palo Alto PA-5060	Cisco ASA 5585	Validated By
On Device Management	Supported	Supported	Supported	
Centralized Management	Provider-1	Panorama	Cisco Security Manager	
Routing Management	Provider-1	Panorama	Cisco Works	
Firewall Management	Provider-1	Panorama	Cisco Security Manager	
NAT Management	Provider-1	Panorama	Cisco Security Manager	
VPN Management	Provider-1	Panorama	Cisco Security Manager	
Content Filtering Management	Additional Addon Software	Panorama	N/A	
Antivirus Management	Additional Addon Software	Panorama	N/A	
IPS/IDS Management	Additional Addon Software	Panorama	Cisco Security Manager	

Virtualization Management	Additional Addon Software	Panorama	N/A		
HA Management	Provider-1	Panorama	Secure Device Manager		
Score	5	4	2		

- Evaluated in Researcher's Lab
- Evaluated by NSS/Gartner

In the past, many of the functions we have evaluated would be found in differing equipment which would drive different management touch points. One of the main benefits of UTM is the ability to manage many of the unified threat techniques from a common “pane of glass”. Table 4.15 shows what criteria were evaluated in the area of unified management and how each vendor implemented the management of the technology.

The first evaluation point is whether the device can be managed locally, centralized or both. Most devices that act as a standalone system can be managed locally but as the device counts start to grow, it is valuable to have centralized management. All three products supported this model. The remaining features of the UTM appliance are subdivided to indicate what software package manages each. Check Point has a strong history with Provider-1 which is the unified management interface for their product. It manages all aspects of the UTM model with some additional software add-ons. Palo Alto has the Panorama software package which does similar functions as Provider-1. Cisco, with its lack of features in the traditional sense, obviously has some gaps with centralized management.


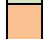
Provider-1 has a proven track record in the industry and the ability to navigate and effect change in an intuitive way is primarily noted. Panorama is certainly a challenger in this space but given the market maturity of the Provider-1 product, Check Point edge out the competition in this area.

Unified Logging

Logging of information is the output of the security perimeter and the overall status of it. Many times it is intended to alert or inform the security team that something is outside the scope of normal. Logging has been present since the beginning of all of the evaluated platforms but similar to the way that management has been unified, logging of the various devices has also been unified. To what degree is what our research quantified. At the low end of the scale, unified logging could just mean individual components in the same box are now all placing logs into one location.

Table 4.16 Operations – Unified Logging

Operations - Unified Logging							
	Check Point Power-1	Palo Alto PA-5060	Cisco ASA 5585	Validated By			
On Device Logging	Supported	Supported	Supported				
Centralized Logging	Supported	Supported	Supported with Software Addon				
Syslog Compatible	Supported with Software Addon	Supported	Supported				
Open Standards Log Format	Not Supported	Supported	Supported				
Unified Logging of All Events	Supported for Capabilities of Platform	Supported for Capabilities of Platform	Supported for Capabilities of Platform				
Unified Reporting	Supported for Capabilities of Platform	Supported for Capabilities of Platform	Supported for Capabilities of Platform				
Exportable Logs	Supported	Supported	Supported				
Score	5	5	2				

 Evaluated in Researcher's Lab
 Evaluated by NSS/Gartner

At the high end, the intelligence of having all of the functions logging in a single platform could allow the system to better correlate what is actually happening and thus provide

more information to security professionals, saving time on tracking down various pieces of information. Table 4.16, similar to the previous table, indicates that each vendor provides both on-device and centralized logging, although Cisco's is a software add-on. One area of obvious concern immediately is that the Check Point logging is a proprietary format. It has been this way for some time though and tools are available to convert the Check Point logs into standard syslog format. Both Check Point and Palo Alto supported unified event logging for the UTM features they support and both excelled in this area.

Exporting the logs into other formats is a strong desire but the ability to see real time logs on the device during troubleshooting provides a valuable asset. In both Check Point and Palo Alto platforms, they have a rotational logging structure that allows for fast access to the logs locally. Also the ability to execute real time logging with tools such as TCPDUMP is in both platforms. Because both platforms allow the exporting of log data to syslog outside tools can be used to draw correlations. Because the purpose of each device is the UTM functions, we rely on other vendors who excel at taking in this information and drawing conclusions. In this case, Check Point and Palo Alto both make it more than easy to accomplish this and thus the scoring in this area was equal.


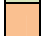
Command Line Interface

Before there were graphical user interfaces into the management of these platforms, command line input was the popular way of configuring and operating these devices. Although GUIs have picked up in dominance of usage for management, CLI is still used often by people who are comfortable with them and also in situations where there are parameters that can only be changed via the CLI. Granular control at the CLI is still vital to the UTM platforms, so having a well architected, intuitive and easily navigated CLI was worth evaluation. One of the most

popular CLIs in the industry is the Cisco CLI. The familiarity with the CLI simply allows operational staff to pull from previous experiences to extend their ability to support the platform with less education needed. Because Palo Alto uses a Cisco-like CLI, we scored Cisco and Palo Alto higher in this category as shown in table 4.17.

Table 4.17 Operations – Command Line Interface

Operations - Command Line Interface							
	Check Point Power-1	Palo Alto PA-5060	Cisco ASA 5585	Validated By			
CLI Access	Supported	Supported	Supported				
CLI Type	IPSO - Unix Like	Built on BSD with Cisco-Like Command Structures	Cisco CLI				
Score	3	4	4				

-  Evaluated in Researcher's Lab
-  Evaluated by NSS/Gartner



Policy Conversion

Policies are the rules, configurations and parameters that are set inside of an appliance that instruct the security device on what to do. In traditional firewalls, policies were the rules that outlined source, destination and ports that were the criteria for permitting or denying traffic. Some of the traditional firewall policies are thousands of lines long which explicitly identify certain types of traffic. One of the concerns of large enterprises is that these lists of rules would have to be recreated inside of any new technology. Because of this, evaluation of how easy it is to import policy into the new UTM devices was something that was worth researching. Each vendor has software packages that allow for policy conversion between various platforms. In the case of Palo Alto and Cisco, they natively support conversion of policy from Cisco and Checkpoint with Cisco allowing for Netscreen/Juniper conversion as well. Check Point supports

Cisco and Netscreen/Juniper but only with a separate software package. With native support, the scoring shows that Palo Alto and Cisco scored higher in this area. The native tools are straight forward and Palo Alto was able to convert a complex policy of thousands of lines in a matter of days. Table 4.18 displays the scoring for this area of the operations evaluation.

Table 4.18 Operations – Policy Conversion

Operations - Policy Conversion							
	Check Point Power-1	Palo Alto PA-5060	Cisco ASA 5585	Validated By			
Policy Conversion Tool	FirePac - Separate Software	Native Converter for Cisco and Check Point	Native Converter for Check Point and Netscreen				
Score	3	4	4				

-  Evaluated in Researcher's Lab
-  Evaluated by NSS/Gartner

Support and Cost

Education is also a consideration when selecting a platform. How available are the classes to become educated on the vendor’s equipment? Are there certifications for becoming an expert on a vendor’s technology? Is the equipment widely deployed enough where there may be a lot of information in online forums or white papers that could offer more insight? These are all questions that should be taken into account with respect to support.

When buying any technology, one of the evaluation criteria is the support that a buyer can expect to receive from the vendor. Reputation sometimes can play an important factor in this as some people acknowledge that larger vendors will be better staffed to handle the support needs. Others consider that smaller more innovative companies are more amenable to personalize the support and allow for customized implementations specific to the company.

Support comes in terms of reactive assistance when something is not working on the device but can also refer to proactive information such as product improvements, long term road maps and vision.

Lastly, the cost of the platform is going to be a key differentiator. Costs are hard to quantify because of the options that an appliance can be configured with in addition to whether you must buy the equipment through a reseller or direct from the vendor. Support costs can vary as well depending on how difficult the platform is to support for the vendor. Because there are so many parameters that can affect the Capex and Opex of the solution, only a general and terse look at the expense of these platforms was done. The cost section is only intended to provide color to the more important areas discussed above.


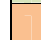
Education

Getting up to speed on new technologies can always be arduous especially with a new player in a technology. If the vendor is a new challenger into the field, it is possible they do not have formalized training opportunities to become familiar with the product. In this case organizations would have to rely on the vendor for customized in-house training, which could be a benefit depending on how structured the training would be. Lab testing is always a good way to take a new platform for a road test so the ability for equipment to be loaned or demonstrated before buying is also a consideration. With most technology platforms, becoming a master in that platform can provide the interested companies with certification benchmarks that indicate how proficient someone is in that technology. This is extremely helpful in staffing. It can also in some cases allow for lower support costs from the vendor as they acknowledge the qualified staff that is on hand. Market dominance ultimately comes into play in this area. While all three vendors offer very competitive classes for the equipment they produce, market maturity would

dictate that Check Point and Cisco would be much further along than Palo Alto which is clearly identifiable in table 4.19. This is also indicative of the number of different technologies each develops for. The two leaders in this space also have several industry aged certifications. Palo Alto, which is still emerging in this space, is still working on a certification program for their equipment.

Table 4.19 Support and Cost – Education

Support and Cost - Education							
	Check Point Power-1	Palo Alto PA-5060	Cisco ASA 5585	Validated By			
Available Public Classes	Over 230 Partners Offer Check Point Classes	Moderate Number of Partners Offer Palo Alto Classes	Over 500 Partners Offer Cisco Classes				
Available Certifications	CCSA, CCSE, CCMA	Not Yet Available	CCENT, CCSP, CCIE, CCNA Security, CCNP Security				
Custom Training	Supported	Supported	Supported				
Score	4	3	4				

 Evaluated in Researcher's Lab
 Evaluated by NSS/Gartner

Support

As mentioned above, support is a key component of the purchase of any equipment. In the case of UTM, with so many functions coming together it would be logical that support for these devices needs to be comprehensive and smooth. The first level of support is the support team which consists of the account manager and the sales engineers. Together this team should be well immersed in the goals and objectives of the perimeter. They should be experts in their equipment and knowing how it provides the best possible solution for a given set of

requirements. A technical assistance center should also be available 7x24x365 in order to assist with any failures or problems incurred with the equipment. While reactive support is an absolute must, proactive support is also something that should get consideration. It is important to stay tuned into the vendor and how they are evolving their platform to meet future demands. There should be a two-way dialogue between the customer and vendor where needs and requirements from the customer are funneled back to the vendor for incorporation into the equipment road map. The vendor should have a plan, a vision for where they see the market going and how their equipment intends to provide value in that direction.


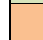
For the testing, each vendor’s support model was evaluated. In each case, local support engineers are available in most major locations in the U.S. Web site complexity and the ability to navigate for FAQs, help files, software downloads and general information showed that the two market leaders have a bit more complexity. This is again indicative of the sheer amount of product they support. It is certainly easier to stay simple when the product portfolio is small.

Table 4.20 shows the complexity of getting access to a technical assistance center or TAC. This

Table 4.20 Support and Cost – Support

Support and Cost – Support

	Check Point Power-1	Palo Alto PA-5060	Cisco ASA 5585	Validated By
Local Support Available	Supported	Supported	Supported	
Complexity of Website	Moderate	Simple	Moderate	
Complexity of TAC	Heavy	Moderate	Heavy	
Access to Developers	Difficult	Simple	Difficult	
Access to Road Map Information	Simple	Simple	Simple	
Score	2	4	3	

 Evaluated in Researcher's Lab
 Evaluated by NSS/Gartner



area mirrors that of the aforementioned as does the access to developers. In the testing of the Palo Alto, developers were on hand to assist in the evaluation, explain how the platform operates and offer any changes to the different UTM modules. Palo Alto as an up and comer is obviously fighting hard for business and given their market focus, their support was very personalized and attentive which resulted in a higher score than the other two vendors.

Cost

Because the purpose of the thesis was to show support for how traditional security devices must evolve into a UTM model and the focus was on the technological reasons, costs were only an addition to provide some perspective in summation of the other elements. Cost of the device can be a difficult thing to quantify because one size does not fit all when it comes to these UTM platforms. For research purposes, we outline costs of the evaluated platforms to attempt to show comparisons from the aforementioned benefits to a cost ratio. Figure 4.21 shows the relative costs for each platform as it relates to what was evaluated.

Table 4.21 Support and Cost – Cost

Support and Cost – Cost							
	Check Point Power-1	Palo Alto PA-5060	Cisco ASA 5585	Validated By			
Cost - Standard Chassis	~\$64,000	~\$40,000	~\$70,000				
Cost - Fully Loaded (All Features)	~\$200,00	~\$150,00	~\$115,00				
Cost – Support			~\$14,000				
Score	2	4	3				

	Evaluated in Researcher's Lab
	Evaluated by NSS/Gartner

Chapter 5 – Conclusions

Network perimeters are under attack by new threats that seem to be launching ever so quickly. Threats are aimed at the network, at the host and application but the primary exposure is to the organization's data. Data is the true asset of the company. During an ever increasing time of threats, business and markets are pushing toward even more communication inside and outside of the secured perimeter. Over the past 25 years, risk management has attempted to solidify the security model around these competing requirements. Too little security allows for flexibility but exposes vulnerabilities more. Too much security stifles the organization by suffocating the access. The result of this effort has been the consistent deterioration of the effectiveness of firewall technology.

The first conclusion to be drawn is that a methodical approach to security management should be followed. Without the processes and procedures discussed in the early sections of the thesis, the technology will not be deployed, configured or operated in any efficient way. Without security policies that have been well analyzed, the products will not meet the objectives. Misidentification of a risk could be an end-game mistake.

Unified threat management is an approach to consolidate many of the tools that are used in mitigating these risks. It consists of the combination of disparate technologies today into a single core platform. From the research above, the most obvious observation is that the industry sees a real opportunity to consolidate down the number of devices that exist in the perimeter network. This obviously shrinks complexity, operational costs and creates a more efficient packet flow. It is however grounded with some very real concerns. Some organizations are concerned about putting all of the functionality into "one basket". If a single vendor controls many of the mitigation techniques that were separated before, that vendor is now on the center

stage to handle all of these. From test results, it was apparent that no device has the ability currently to have all features turned on and still perform at the speeds listed. One feature could affect the processing of another. As hardware becomes more mature, this can change but the important take-away is that the unified direction is set. So one must weigh the potential benefits with the risk factor associated with that consolidation.

For that reason, the market shows the most mature area of the research is the next generation firewall. The basic idea of the firewall was to permit or deny traffic into and out of the enterprise network. Over the years, this device fell behind the advances in applications and exploits of them. The NGFW is the core strategy that should be recognized from the results of this thesis.

New concepts being introduced in the NGFW such as application identification, single pass technology for increased packet efficiency and the addition of other technologies such as IDS/IPS, content filtering and antivirus are the basis for UTM but it should be recognized that vendors such as Palo Alto and Check Point have centered their focus around perfecting the NGFW. Figure 2.16 in the research depicts that the NGFW is the heart of the UTM effort. If the core of the platform is not able to change to accommodate the flaws of the past, then vendors are simply throwing technologies together in the same chassis with no real innovation. Vendors need to perfect this core to the extent that applications are inspected regardless of port and do so at wire speed. Once the application is inspected, and the packet is open, then apply all of the different technologies in an accelerated process to expedite the forwarding. Functionally the NGFW seems to be the strongest movement for UTM. UTM is still immature in the market as shown by much of the matrix but the firewall features that form the NGFW seem to be maturing at a much quicker rate. According to Gartner research (2010), next generation firewalls account

for only 1% of the Internet connection security mechanisms today. They believe that by 2014 that number will be increased to 35% install base and that 60% of all new purchases will be NGFW. The predicted rise of nearly 34% in three short years reinforces the research data in this thesis. The current perimeter technologies are not meeting the challenges today or in the future with respect to risk.

The results support the notion that while it is possible to place all of the functions into a single chassis, it is not a perfect model. Performance will vary as more and more things are turned on. Single pass technology shows promise for getting the inspection needed at hardware accelerated speeds. Operations are certainly on their way to being streamlined as more things are consolidated. There is still a long way to go with providing a “single pane of glass” view into the enterprise security but the efforts made so far have shown promise that vendors realize that operating environments from a support perspective must change. As device counts are cut down in the gateway, the overall costs of staff, equipment, support and environmental should begin to decrease.

This thesis has attempted to provide solid footing to the UTM effort. Not to say that UTM as a concept is ready for the market but that through focused energy on the NGFW, perimeter security inches closer to the idea behind UTM. Future researchers have the opportunity to take the research further by examining challengers to the NGFW quadrant. Because secondary research has shown that small companies with lower bandwidth requirements and less complex environments are more apt to deploy UTM, it would be valuable to research the penetration of UTM in the small to mid-sized business sector. Comparatively it would be ideal to show how large enterprises are gravitating more towards the NGFW concept and to theorize where the two concepts will start to blend.

UTM is not necessarily a product today but more an idea that the evaluated vendors are working toward. So far the solutions evaluated show the course is set and with Gartner recognizing this space in their future quadrants the only speculation is that this will be a growth sector. For now, the platforms focus on a core design consisting of a NGFW with fully integrated threat protection that runs on customized hardware giving it the ability to meet security with the performance requirements.

Chapter 6 – References

- Al-Radhi, A. (Producer). (2009, 7-29-2011). Network security 2.0. [PowerPoint Slides]
Retrieved from <http://www.menog.net/menog-meetings/menog4/presentations/Network%20Security%20V2.0.pdf>
- Charette, R. (1996). *The risk with risk identification*. ITABHI. Retrieved from
http://www.itmpi.org/assets/base/images/itmpi/privaterooms/robertcharette/RISK_ID.pdf
- Check Point Software Technologies. (1994). *Check Point introduces revolutionary Internet firewall product providing full internet connectivity with security; wins 'best of show' award at Networld+Interop '94*. Retrieved from
http://www.checkpoint.com/press/1994/interop_press.html
- Cisco Systems. (2009). *Cisco Visual Networking Index: Forecast and Methodology, 2009-2014* [Whitepaper]. Retrieved from
http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html
- Coffman, K. & Odlyzko, A. (2001). Internet growth: Is there a "Moore's Law" for data traffic?. Retrieved from <http://www.telecomvisions.com/articles/pdf/att.pdf>
- Currier, G. (2011). Exclusive Research: Enterprise Security Spending Trends. *CIO Insight, January/February 2011*. Retrieved from http://circ.ziffdavisenterprise.com/creative-services/pdfs/CIOI_2011_0102_Research.pdf
- Doctor, B., & Poynter, I. (2003). Beyond the firewall: The next level of network security. Retrieved from http://www.stillsecure.com/docs/StillSecure_BeyondtheFirewall.pdf
- Forrest, S., & Ingham, K. (2002). A history and survey of network firewalls. Retrieved from <http://www.cs.unm.edu/~moore/tr/02-12/firewall.pdf>

Fortinet. (2011). Accelerating UTM Specialized Hardware. Retrieved from

http://www.fortinet.com/doc/whitepaper/Accelerating_UTM_Specialized_Hardware.pdf

Gosal, S. (2006). Unified Threat Management. *IT Security*. Retrieved from

<http://www.itsecurity.com/features/unified-threat-management/>

Greene, T. (2007). Next-generation firewalls will need wide variety of features. *Network World*.

Retrieved from <http://www.networkworld.com/news/2007/091407-next-generation-firewalls.html>

Hubbard, D. (2009). The failure of risk management: Why it's broken and how to fix it

Retrieved from

<http://books.google.com/books?id=u2AceU1L95EC&lpg=PP1&pg=PR4#v=onepage&q&f=false>

Innella, P. (2001). The evolution of intrusion detection systems. Retrieved from

<http://www.symantec.com/connect/articles/evolution-intrusion-detection-systems>

Kouns, J., & Minoli, D. (2009). Information Technology Risk Management in Enterprise

Environments: A Review of Industry Best Practices and a Practical Guide to Risk Management Teams Retrieved from

http://books.google.com/books?id=UOyGnlbho8wC&pg=PA92&dq=ISO+31000&hl=en&ei=K9xWTpTTLYOQsQLxvPWhDA&sa=X&oi=book_result&ct=result&resnum=2&ved=0CD0Q6AEwAQ#v=onepage&q=ISO%2031000&f=false

Menninger, M. (n.d.). The story of the first internet worm. *Buzzle.com*. Retrieved from

<http://www.buzzle.com/editorials/10-10-2005-78536.asp>

Messmer, E. (2010). What you should know about next generation firewalls. *Network World*.

Retrieved from <http://news.idg.no/cw/art.cfm?id=A40770F4-1A64-6A71->

[CEC4760562664D1A](http://news.idg.no/cw/art.cfm?id=A40770F4-1A64-6A71-CEC4760562664D1A)

NSS Labs. (2011). Next-Generation Firewall Individual Product Test Results. Retrieved from

<http://www.checkpoint.com/campaigns/nss-next-gen-firewall/index.html>

risk. (2011). In *OxfordDictionaries.com*.

Retrieved from <http://oxforddictionaries.com/definition/risk?view=uk>

Northcutt, S., Zeltser, L., Winters, S., Frederick, K. K., & Ritchey, R. W. (2003). *Inside Network*

Perimeter Security. Indianapolis: New Riders Publishing.

Smaha, S. E. (1988). Haystack: An intrusion detection system. *Proceedings of the IEEE 4th*

Aerospace Computer Security Applications Conference (pp. 37-44)

Wai, L. (2001). Security Life Cycle - 1. DIY Assessment. *Sans Security*. Retrieved from

http://www.sans.org/reading_room/whitepapers/testing/security-life-cycle-1-diy-assessment_260

Weaver, R. (2007). *Guide to Network Defenses and Countermeasures*. Boston: Thomson.

Young, G., & Pescatore, J. (2010). Magic Quadrant for Enterprise Network Firewalls. Retrieved

from <http://www.paloaltonetworks.com/cam/gartnerMQ/register.php>

Glossary of Terms

ASIC – Acronym for ‘application specific integrated circuits’. It is a chip that is designed for a specific application rather than a generic microprocessor.

BGP – Acronym for ‘border gateway protocol’. It is an exterior dynamic protocol that is used to commonly connect differing autonomous systems together.

CLI – Acronym for ‘command line interface’. It is the visual interface that allows a user to interact with a devices operating system.

DSL – Acronym for ‘digital subscriber line’. It is a broadband technology offered by telecommunications companies to connect to the Internet.

EMS – Acronym for ‘element management system’. It is an application that allows for management of network elements in a centralized manner.

FAQ – Acronym for ‘frequently asked questions’. It is a list of questions that are most commonly asked with answers provided.

Firewalls – A device that is used to inspect and filter traffic on a data network. It uses policies and rules to determine what traffic is permitted and what is denied.

GUI – Acronym for ‘graphical user interface’. It is an interface that is used to allow humans to visually interact with a computer’s operating system.

H.323 – It is a standard protocol that is used to provide audio and video communications on data networks. It is a signaling protocol that provides for setup and teardown of a session.

HA – Acronym for ‘high availability’. It is the concept of providing redundancy into an environment by adding active mirrors of devices into the traffic flow whereas in the event of a failure other devices are able to actively take over.

HIPAA – Acronym for ‘health insurance portability and accountability act’. This national standard provides protection for health patients to protect their personal information.

HTTP – Acronym for ‘hypertext transfer protocol’. This is the primary protocol that constructs the World Wide Web and allows for users to connect to web pages.

IDS – Acronym for ‘intrusion detection system’. This device is responsible for monitoring traffic and identifying when an intrusion is likely happening.

IP – Acronym for ‘internet protocol’. This is the primary protocol that allows computing devices to communicate with each other at the network layer.

IPS – Acronym for ‘intrusion prevention system’. This device is responsible for monitoring traffic and not only identifying when an intrusion is occurring but also preventing such intrusion.

LDAP – Acronym for ‘lightweight directory access protocol’. It is an application protocol used for querying and controlling directory services which can provide authentication control for an enterprise.

MAC – Acronym for ‘media access control’. It is the layer of the network that contains a hardware based address that is uniquely identified to a specific vendor. MAC addresses are normally contained at the data link layer.

NAT – Acronym for ‘network address translation’. It is the process of changing the source, destination and/or ports for a given communication path.

NGFW – Acronym for ‘next generation firewall’. It is the term used to describe the emerging firewalls that contain new features such as application identification, high speed packet inspection and elements of unified threat management.

OSI – Acronym for ‘open systems interconnection’. It is the term used to describe the framework for how, using a layered approach, communications between two end points should be represented.

OSPF – Acronym for ‘open shortest path first’. It is an interior dynamic routing protocol that is commonly used inside of enterprise networks for the distribution of routes.

PCI – Acronym for ‘payment card industry’. It is a term used to describe the process of securing any transaction that contains sensitive payment card information.

QoS – Acronym for ‘quality of service’. It is a term used to describe the methodology of identifying key traffic types and providing a level of service appropriate for that traffic type.

RIP – Acronym for ‘routing information protocol’. It is an interior dynamic routing protocol that is commonly used inside of enterprise networks for the distribution of routes.

Router – A device that is responsible for guiding packets along through an interconnected system. It utilizes packet information to decide where traffic should be sent.

RTP – Acronym for ‘real time protocol’. It is a protocol for providing transport for real time applications such as voice and video.

SCCP – Acronym for ‘skinny client control protocol’. It is a Cisco proprietary protocol that is used for signaling a call between two endpoints on either voice or video.

SIP – Acronym for ‘session initiation protocol’. It is an industry standard protocol that is used for signaling a call between two endpoints on either voice or video.

SOX – Acronym for ‘Sarbanes Oxley’. It is legislation that dictates which business records must be retained and for what period of time.

SRTP – Acronym for ‘secure real time protocol’. It is the secured version of RTP.

TAC – Acronym for ‘technical assistance center’. This is the vendor supplied center that a customer would call in order to report problems with a device, software or service.

TCP – Acronym for ‘transmission control protocol’. A protocol at layer 4 of the OSI model that is responsible for a connection based communication path that involves setup, flow control and teardown.

TOS – Acronym for ‘type of service’. It is a field in the IPv4 header that has been traditionally used to mark packets for quality of service treatment.

UDP – Acronym for ‘user datagram protocol’. Similar to TCP in that it operates at layer 4 of the OSI model. UDP is responsible for packet delivery but is not connection oriented and does not have any delivery guarantee.

UTM – Acronym for ‘unified threat management’. A term used to describe the consolidation of threat management techniques into a more cohesive platform or arrangement.

VLAN – Acronym for ‘virtual local area network’. A technology that allows for broadcast domains to be logically spread across physical devices.

VPN – Acronym for ‘virtual private network’. It refers to a private network that configured within or using a network that is not controlled by the private network owners. It allows for the extension of a private network across uncontrolled boundaries.