Network-Based Security Requires Firewall, IPS and AV

White Paper

High Performance Multi-Threat Security Solutions

FORTINET™

Introduction

When it comes to information security these days, it's a mixed up muddled up world out there. The terms being used to describe network defense capabilities are just as blurry and hard to pin down as the latest flavor of blended threat. Not surprisingly, the result is a growing state of misunderstanding and confusion, culminating in the inability of organizations to readily separate fact from fiction. Indeed, amidst the haze of imprecision there is even a proposition that achieving comprehensive network-based protection requires little more than intrusion prevention and, of course, firewall capabilities.

The intent of this paper is to set the record straight, both in terms of clearing up some of the related imprecision, as well as by definitively demonstrating that antivirus (including anti-spyware) is also an essential, complementary component of an organization's network-based defenses. Indeed, it takes all three – robust firewall, intrusion prevention, and antivirus capabilities – to form the core of an ideal network-based security solution.

Building a Solid Foundation

One of the greatest challenges facing organizations today is the inconsistency and imprecision that plagues the terms used to describe the threats that are being encountered and the countermeasures that are being sold to address them. For example: what does "application layer" really mean? What is the difference between a virus, a worm, and malware? And, just as important, does it really matter?

Accordingly, it is essential to establish a common language before launching into a discussion of how the threat-scape is evolving and the impact this is having on the security solutions that organizations should be implementing.

Layers Can Be Confusing

The qualifier "application" or, more commonly "application layer", is quite possibly the most confusing and, depending on the perpetrator, misleading term in information security today. For example, what exactly is an "application" threat vector? Or what does it mean when a vendor says their product has extensive "application layer" protection capabilities?

The issue here is two-fold. The first problem is that there is no single, definitive taxonomy that has been uniformly adopted by all parties in the security community, vendors and users alike. Consequently some vendors will establish their own definitions for key descriptors and terms, while others will purposely leave them vague, if not couched in generous references, to encourage advantageous interpretations. The second problem is that the one, essentially de-facto framework that is relatively well known, the Open Systems Internetworking (OSI) reference model, is often misunderstood – and again, purposely misused.

In this regard, it is important to realize that the OSI 7-layer model is actually intended to describe a modular approach for communications between networked end users (see *Figure 1*). In other words, it is fundamentally a *network* communications model and therefore has very little to do with actual *applications* (e.g., a web browser, email, SAP) – other than in terms of conveying their content and commands as payload. Even the unfortunately named "application layer", where protocols such as SMTP (for email) and HTTP (for web) reside, is really just about providing services to ensure that higher-order applications can communicate across all types of environments. To complicate matters even further, there are in fact some "network applications" that reside within OSI Layer 7 (i.e., they can be utilized independently of a higher-order application). Telnet and FTP are two examples. However, their presence only serves to further validate the point that the term "application layer" can easily be misrepresented and/or misunderstood.

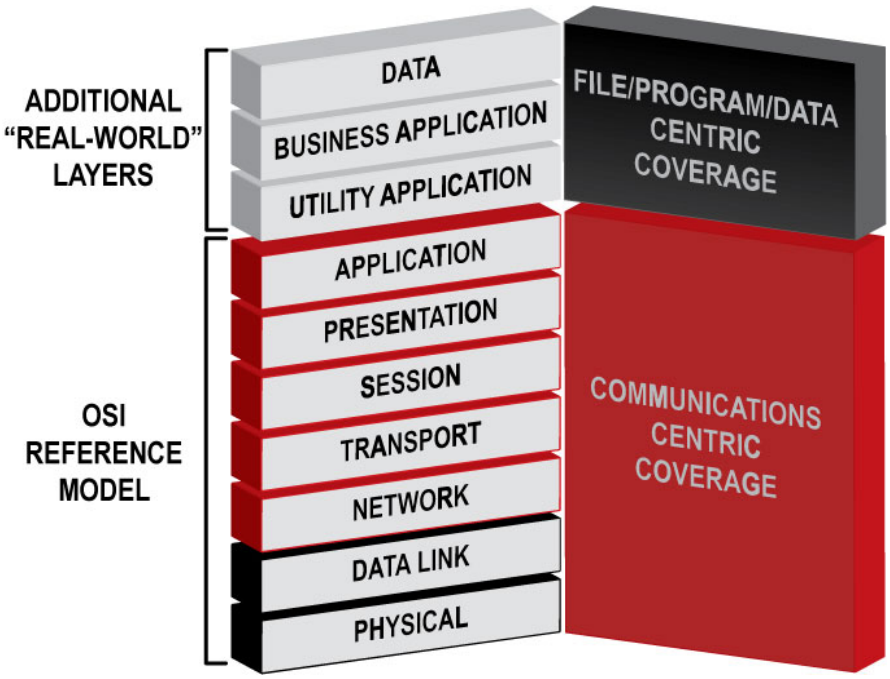

Figure 1: The OSI Reference Model and Beyond

What does all of this mean when it comes to information security? Well, the key points most relevant to the topic at hand are as follows:

- Care should be taken when interpreting the meaning of phrases such as "application layer threat" (or even just "application threat") and "application layer protection", particularly when they are used in marketing materials. There is no definitive or universally accepted convention in use.
- While the OSI "application layer" is indeed distinct from the network and transport layers, they are all focused on network communications. Consequently, there are still higher layers that will require protection as well. For instance, attacks could be focused on compromising: the code/commands of various utility applications (e.g., browsers, web servers, databases), the code/logic of pre-packaged or custom-built business applications (e.g., Word, SAP), or even individual data elements (e.g., personal health information, SSNs).
- A potentially more accurate or at least helpful way to divide the problem/solution space is to consider: communications oriented threats/countermeasures, file oriented threats/countermeasures (i.e., those associated with executable programs or documents with executable content), and data oriented threats/countermeasures (treatment of which is beyond the scope of this paper).

Location, Location, Location

Unfortunately, the use and interpretation of "network" as a qualifier also presents some challenges. In particular, strictly speaking "network security" should be about protecting the network itself – that is, the devices that comprise it, such as routers and switches, as well as the protocols they use to communicate. However, the term is often used instead to refer to any security measure that is on/within the network. As a result, lacking any contextual clues, or worse, faced with ones that are misleading, the meaning of "network security" and the capabilities of any product described using this term can be unclear.

To help resolve this ambiguity, "network-based security" is used henceforth to refer to security measures that reside in the network, as opposed to ones that reside on hosts such as desktops and servers. That said, it is important to realize that this term is only meant to convey a deployment location. It has nothing to do with the scope of capabilities that an associated product may have. In particular, network-based security solutions are not necessarily limited to addressing network layer threats. They can just as easily incorporate countermeasures that are focused on any or all of the following: network layer communications, application layer communications, program and document files, or individual pieces of data.

To reiterate, "network-based" is meant only to refer to a product's location, not to the specific capabilities that it includes. In terms of general capabilities, it is important to acknowledge one significant advantage of a network-based solution over an equivalent host-based countermeasure. By virtue of its location, the network-based option enables adherence to one of the guiding principles of threat protection, which is to stop threats sooner rather than later – in this case, at a network perimeter or chokepoint, ideally before they have spread among the host population.

The Lifecycle of a Threat

Instead of focusing on semantic issues, the purpose of this section is to provide a clear definition of how threats work. Much like clarifying the terms "application layer" and "network-based", this will prove helpful when discussing changes in the threat landscape and the impact they are having on obtaining effective security solutions.

In general, threats consist of the following four stages:

- **Transmission** refers to the process of how a threat gets from its source to its target. For locally executed threats this is accomplished by manually conveying and loading an infected file or otherwise entering code directly into the target system. For remotely executed threats this stage depends on some form of electronic communication and associated protocols (e.g., IRC, instant messaging, file shares, email/smtp, or even just underlying TCP/IP or UDP/IP sessions).

- **Penetration** refers to how a threat subsequently gets into/on its target. For locally executed threats this stage is not really separable from transmission. It involves bypassing system controls or simply taking advantage of permissions granted to local users. In contrast, the remote execution case involves taking advantage of a vulnerability – either in terms of a mis-configuration or a weakness in a piece of code that the target runs. Significantly, such vulnerabilities can be associated with any of the system's communications services (i.e., at any of the OSI layers), or any of the higher-order applications that it runs (i.e., file/program-centric).

- **Launch** refers to the execution of a threat's payload. It is the primary stage associated with doing harm, such as stealing information, over-writing or manipulating stored data, or crashing a process. It may also involve a communications component that enables the attacker to continue to download and execute additional payloads over time.

- **Propagation** is an optional but very common stage that involves perpetuation of the threat, i.e., reproduction and spreading. For example, a threat may scan an address range looking for other hosts, or email itself to all of the addresses in a user's email program. In some instances, self-propagation may be the only "payload" that a threat effectively has – though as will be discussed shortly, this semi-benign mode of operation is rapidly becoming less common.

Figure 2: Threat Lifecycle

Another purpose for outlining the lifecycle of a threat is to highlight the fact that there are many points at which an attack can be thwarted. As will be discussed shortly, different countermeasures tend not only to focus on different "layers" of the problem, but also on different stages of the problem. As a general rule, stopping a threat earlier in its progression through the above stages is preferable since it will typically provide the greatest reduction in terms of the threat's impact on the target environment.

The Threat-scape

Understanding the nature of threats and how they are evolving is the next prerequisite for establishing the components and requirements of a solution that is capable of countering them.

What's In a Name?

Unfortunately, the chore of sorting out the meaning of associated terminology is not quite finished. One piece of the puzzle that still remains is the names used to describe various types of threats. After having been used so frequently in recent years, often inappropriately, they have now lost much of their original distinction. To help remedy this situation, brief attribute-based descriptions are provided below for some of the most relevant threat types.

Virus

- A program that is able to self-replicate
- Depends on user interaction for penetration, execution, and propagation
- Depends on a host file/program (i.e., is not self-contained) and host resources (e.g., disk)

Worm

- A program that is able to self-replicate
- Takes advantage of a vulnerability to penetrate and execute without any user interaction
- Self-contained in that it resides in system memory and uses system and network services for transmission and propagation – it has no reliance on files/programs.

Trojan

- A malicious program that is within or masquerading as an innocent or useful program
- Does not self-replicate
- Does not infect other files

Of course, these only represent a portion of the overall threat pantheon, which also includes spyware, targeted attacks, rootkits, bots, phishing, Denial of Service (DOS) and Distributed Denial of Service (DDoS) attacks. A general term that applies to all of these threats, at least to the extent that they rely on malicious software, is malware. The increasing prevalence of this over-arching term is appropriate,

particularly since even the more granular descriptions are really just generalizations. Indeed, the industry at-large has struggled from the outset to maintain a rigorously defensible, technical distinction between viruses and worms. And with the emergence and increasing frequency of blended threats over the past few years, doing so has become an even greater challenge – or a moot point, depending on your perspective. The fact is that by utilizing multiple mechanisms to accomplish any or all of the lifecycle stages, any given blended threat will often incorporate characteristics and capabilities of both viruses and worms. That said, understanding the different ways in which these critters can operate is still an essential prerequisite to being able to effectively stop them.

Trends and Tribulations

Equally essential is understanding how threats are changing. In this regard, one very significant characteristic is the decreasing window of time between the announcement of a new vulnerability and the release of an associated threat capable of exploiting it. Problematic to be sure, this condition is indicative of the need for countermeasures that are more proactive in nature. Undoubtedly, it is also at least partially responsible for the aforementioned emphasis being placed on intrusion prevention systems (IPS) relative to antivirus tools (AV), which are generally classified as reactive in nature. However, a couple of key "clarifications" are in order at this point.

First, the classifications of IPS as proactive and AV as reactive are really just generalizations. Real-world implementations of these technologies typically incorporate both proactive and reactive detection capabilities. For example, most IPS solutions include some measure of Intrusion Detection System (IDS)-like signatures, while all of the leading AV tools now include various heuristic and behavior-based techniques.

A second clarification is that the current need for more proactive countermeasures is largely about identifying a historical deficiency – one that has come to light due in particular to the latest additions to the threat landscape. This should NOT be confused with or interpreted as a situation where all other classes of threats have abated, thereby eliminating the need for other types of countermeasures. After all, threats that exclusively target specific vulnerabilities are not the only ones out there.

Indeed, while current malware trends indicate a decline in the plain-vanilla variety of file-infection viruses, several other types of threats which incorporate virus or virus-like components are actually on the rise. These include:

- **Blended threats**, a significant percentage of which include file/program-centric components. For example, the malware FunLove included a virus and a worm, while Bagle.H of the FunBag family essentially involved a virus in a worm (i.e., as one of its many possible payloads).

- **Trojans** of all sorts, which are often distributed via spam or a worm, but which inherently involve loading and operating a program on the target system. This program can then be used to download additional malware, or to stealthily steal data from the target user/system.

- **Targeted attacks**, which generally trade rapid propagation capabilities (involving communications layer exploits) in favor of greater stealth (typically yielded by focusing on file-centric exploits). One particularly disturbing example is that of spy-phishing. In this type of attack, a highly customized (i.e., targeted) email that includes a trojan (or a link that will result in one being downloaded) is sent to a target audience that is pre-disposed to use certain web sites and services. The trojan then simply monitors web traffic, waiting for particular sites to be accessed, at which point it acquires credentials and other valuable information.

- **Rootkits**, which are programs that once installed provide stealthing capabilities for other pieces of malware (e.g., trojans) in an attempt to make them undetectable by various host-based countermeasures.

A few of the key points to take away from this brief tour of the threat-scape are: that worms are indeed a significant problem; that they are increasingly being used to facilitate the transmission and propagation of "traditional" viruses; but, also that they are not the only means for this to happen. Just clicking on a link in an email or an innocuous looking object on a web page can initiate the download of a piece of file-centric malware.

Now, whether that malware actually gets to the intended target is yet another matter – one that will be determined by the capabilities of the countermeasures that reside between the source of the malware and its intended destination or target.

Network-Based Countermeasures

As discussed earlier, network-based security has the advantage of at least intending to stop threats before they are allowed to spread throughout an organization's entire computing environment. But what specific countermeasures should be part of a network security gateway to ensure this intention comes to fruition? Bits and pieces of the answer to this question have already been introduced, such as those that address the different "layers" at which threats operate, the stages through which they progress, and the other characteristics they are acquiring as they evolve. The intent going forward, however, is to answer this question more holistically. This will be done first by identifying the specific countermeasures that are needed to account for all of the aforementioned "bits and pieces", and subsequently by identifying other essential characteristics required of a full-featured solution.

In general, the necessary network-based countermeasures fall into two classifications: those which are based on a positive model of control, and those which operate with a negative model of control. That said, the following disproportionate coverage of the latter category is not a reflection of its relative strength, but rather a consequence of wanting to thoroughly address the growing misperception that having an IPS somehow obviates the need for AV capabilities.

Positive-Model Technologies

Positive-model countermeasures are those based on identifying communications, programs, and data elements that are known to be good or necessary, and then allowing only them or activities involving them to proceed. An inherent by-product of this approach is the indiscriminate elimination of unknown as well known threats – at least to the extent that they are not being conveyed within legitimate items.

A classic example of this approach, and an essential component of a network-based security solution, is the firewall. Representative products are typically communications-centric, providing control from the network layer to the OSI application layer, but not beyond. That said, they also have limited capabilities at the file/program level, in the form of being able to block material based on file type, as determined by the file extension (e.g., exe, vbs). Overall, their span of visibility affords them the ability (and limitation) to stop threats primarily when they are being transmitted or are propagating.

As powerful as they are, positive-model technologies are hampered by two inter-related challenges. First, enumerating absolutely all traffic that is considered legitimate is a substantial task, if not impossible. The ever-deeper levels of granularity needed to sufficiently "lock things down" can quickly become overwhelming. And secondly, it will always be necessary to allow some traffic to proceed. With this comes the potential for threats to proceed as well, under the guise of or within those communications designated as legitimate. This is the very reason why negative-model countermeasures are also necessary.

Negative-Model Technologies

In contrast to the positive-model approach, negative-model countermeasures operate on the basis of identifying communications, programs, and activities involving data elements that are known to be bad, presumably because they are associated with a threat, and then taking some form of action to stop the associated activity.

Intrusion detection/prevention systems and AV are the classic examples in this case.

IPS/IDS. Modern intrusion control systems incorporate a wide variety of detection techniques. Although the focus is on those that are proactive (e.g., vulnerability based signatures, identification of traffic and protocol anomalies), reactive mechanisms are often included as well (e.g., "traditional" IDS-style pattern matching). The advantage, due in particular to the proactive detection techniques, is the ability to thwart unknown threats, as opposed to just "known threats". This is typically accomplished at or before the penetration stage of the threat lifecycle. However, as powerful and attractive as this may be, it is also important to understand that these systems have a limited span of visibility. Specifically, their detection techniques are very much communications centric. Sure, in contrast to earlier generation, network layer-focused systems they can now boast "application layer protection". But as already discussed, this is not the same thing as directly protecting the actual applications, business logic, and data. The key point is that while IPS is extremely powerful, it is not exactly complete. It has limited visibility into the payload (i.e., files, programs, and data) of the packets it examines.

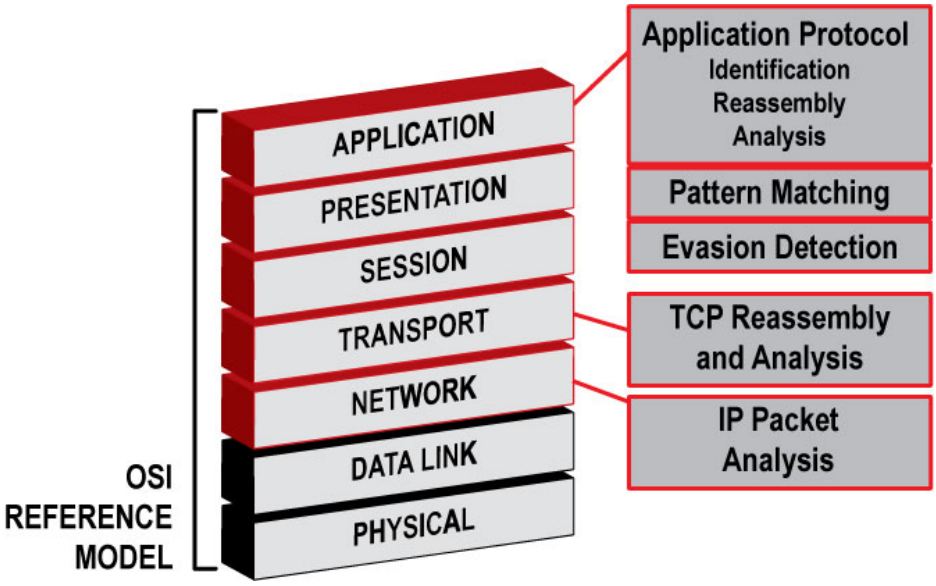

Figure 3: Intrusion prevention techniques and their corresponding Layer

Figure 3 provides a graphical representation of the typical detection techniques associated with IPS/IDS products and the "layers" at which they operate.

AV. In reaction to the realities of the threat-scape, historically reactive AV tools are evolving to also incorporate a handful of proactive detection techniques (e.g., heuristic and behavior-based algorithms). These advancements, however, are far less relevant than the value that AV tools provide in terms of visibility and control at an entirely different layer than IPS/IDS products; while IPS tools are communications centric, AV tools are exclusively file/program oriented. It should also be noted that modern implementations should be inclusive of anti-spyware capabilities. This involves the detection of files, programs or processes that range from outright malware (e.g., password sniffing trojans), to so-called "grayware" (e.g., Active-X plug-ins that track web site usage).

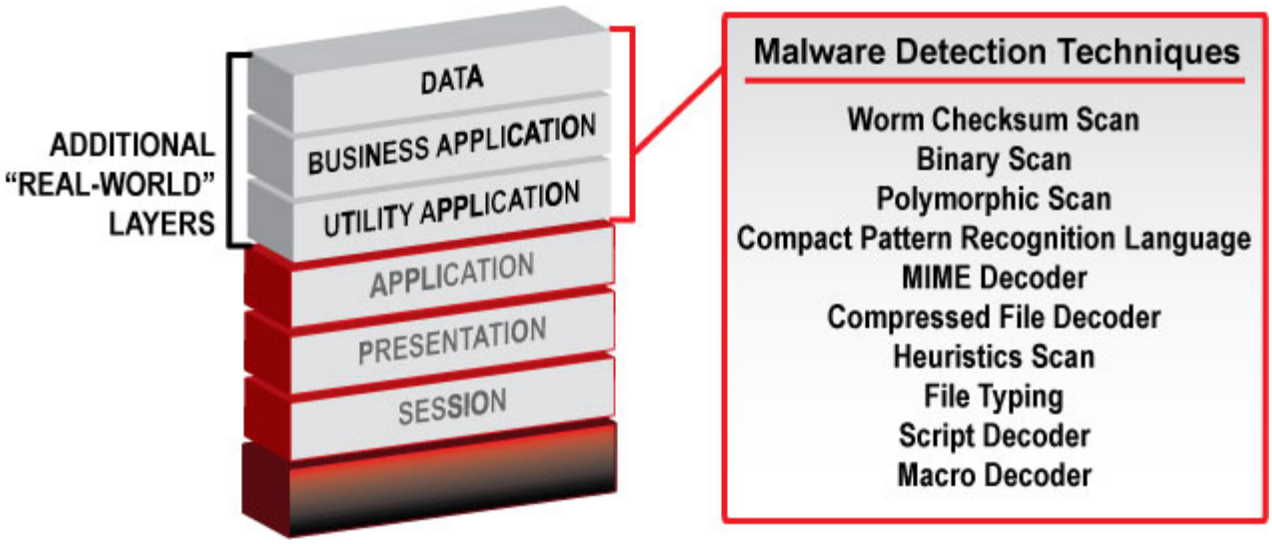

Figure 4: Malware Detection Techniques

Figure 4 provides a graphical representation of the typical detection techniques associated with antivirus/anti-spyware products and their span of visibility.

For the sake of completeness, there is also an extension of the negative-model approach to the data level. Specifically, network-based content filtering solutions can enforce egress/ingress policies based on examination of individual data elements that are included in a communications stream. As such, they have the potential to thwart the execution/launch phase of those threats which involve the harvesting and transmission of sensitive information.

Bringing it All Together

There are two key points to take away from this discussion on network-based countermeasures. First, enumerating all that is "bad" on an a priori basis is no less challenging than doing so for all that is "good". This is what renders positive and negative model approaches complementary, and both types of technologies are therefore necessary for the most effective network-based security solution.

The second point is that, in a similar manner, even within the negative-model category there is a need for complementary technologies. Neither IPS nor AV alone provides the coverage necessary to provide sufficient protection against today's threats. To help crystallize this even further, below are some examples of threat types grouped by the protection technology that is best, and in some cases uniquely, positioned to address representative threats.

Some types of threats that AV controls are uniquely qualified to protect against include:

- Compressed/packed malware (e.g., as with Netsky Bagle and Mytob)
- Polymorphic malware (e.g., as with Uruguay, W95/Marburg, and Magistr)
- Malware embedded in large files – since most IPS tools will only attempt pattern matching on relatively small files (e.g., Viking, Delf.RC)

Some types of threats that network-based IPS controls are uniquely qualified to protect against include:

- Malware and targeted attacks focused on vulnerabilities in communications services (e.g., SQL Slammer, Sasser)
- Denial-of-service attacks (e.g., SYNflood, Teardrop attacks)

Some types of threats where the capabilities of network-based AV and IPS controls overlap include:

- Blended malware, in particular worms that do not exploit vulnerabilities per se, but which are used to distribute user-triggered files/programs (e.g., FunLove, Bagle variants)
- Some types of spyware, particularly those which involve very small files/programs (e.g., Hotbar, ISTBar)

Thus, it should be clear that to be effective a network-based security solution requires, at a minimum, both IPS and AV capabilities, as well as a traditional firewall. In addition, other positive and negative model countermeasures, such as content filtering, VPN, and anti-spam controls, can be used to extend the level of protection even further.

Summary

Once the fog enveloping the terminology used to describe threats and the technologies used to counter them is parted, it should be evident that complementing firewalls only with network intrusion prevention capabilities is not sufficient. Antivirus controls are also a must have. This is validated further by an examination of the evolving threat landscape, which clearly indicates that file and program-centric malware is still quite prevalent – types of malware that intrusion prevention tools are ill-equipped to thwart. Finally, implementing all of these countermeasures (and more) in a multi-layer security platform that is network-based provides the advantage of stopping many threats sooner rather than later – thereby keeping them from spreading more pervasively throughout an enterprise's internal networks and systems.

About the Authors

Mark Bouchard, CISSP, is the founder of Missing Link Security Services, LLC, a consulting firm specializing in information security and risk management strategies. A former META Group analyst, Mark has assessed and projected the business and technology trends pertaining to a wide range of information security topics for nearly 10 years. He is passionate about helping enterprises address their information security challenges. During his career he has assisted hundreds of organizations world-wide with everything from strategic initiatives (e.g., creating 5-year security plans and over-arching security architectures) to tactical decisions involving the justification, selection, acquisition, implementation and operation of their security and privacy solutions.

Freddy Mangum, Vice President of Product Marketing, brings to Fortinet more than 12 years of sales, marketing and business development experience with companies in the networking and security markets. Freddy most recently owned a marketing consulting company that provided product strategy and marketing services to companies such as IronPort Systems, Sarvega (acquired by Intel) and Permeo (acquired by Blue Coat). He was previously employed with prominent security companies, such as Internet Security Systems (ISS), where he directed product marketing activities for product lines generating more than $250 million in revenue. Freddy has also held numerous senior technical marketing and consulting engineer roles with companies such as Cisco Systems, WheelGroup and UUNET.

Chris Simmons, Senior Product Marketing Engineer and the Fortinet Threat Research Team contributed to this document.

About Fortinet

Fortinet is the pioneer and leading provider of ASIC-accelerated multi-threat security systems, which are used by enterprises and service providers to increase their security while reducing total operating costs. Fortinet solutions were built from the ground up to integrate multiple levels of security protection--including firewall, antivirus, intrusion prevention, Web content filtering, VPN, spyware prevention and antispam--providing customers a way to protect multiple threats as well as blended threats. Leveraging a custom ASIC and unified interface, Fortinet solutions offer advanced security functionality that scales from remote office to chassis-based solutions with integrated management and reporting. Fortinet solutions have won multiple awards around the world and are the only security products that are certified eight times over by the ICSA (firewall, antivirus, IPSec, SSL, IDS, client antivirus detection, cleaning and antispyware). Fortinet is privately held and based in Sunnyvale, California.

FORTINET
1090 Kifer Road, Sunnyvale, CA 94086 USA
Tel +1-408-235-7700 Fax +1-408-235-7737
www.fortinet.com

WPR128-0906-R1