



The Power to Prevent:
It's the Next Generation

FIDELIS SECURITY SYSTEMS, INC. 4416 EAST WEST HIGHWAY, SUITE 310, BETHESDA, MD 20814



Mind The Gap: Fidelis XPS™ Deep Session Inspection™

Fidelis Security Systems
Will Irace
Director, Research & Services
December 2010

Version 1.0

INTRODUCTION..... 3

THE THREATS..... 3

THE DEFENSES..... 4

 FIREWALL 4

 INTRUSION DETECTION & PREVENTION 5

 ANTIVIRUS 6

 WEB FILTERING..... 6

 DATA LEAKAGE PREVENTION 7

IT'S NOT WORKING 8

MIND THE GAP: FIDELIS XPS' DEEP SESSION INSPECTION..... 10

 WHY SESSION ANALYSIS MATTERS 10

 HOW IT WORKS..... 1

 THE POWER OF POLICY: THINK FIDELIS! 13

 EXTRAORDINARY POSSIBILITIES 15

 PAYLOAD EXTRACTION AND THE DECODING PATH 1

 INFORMATION FLOW MAP: SEE EVERYTHING..... 17

 OPERATIONALIZING DYNAMIC THREAT INTELLIGENCE 18

CONCLUSION 18

FIDELIS SECURITY SYSTEMS, FIDELIS EXTRUSION PREVENTION SYSTEM, FIDELIS XPS, DEEP SESSION INSPECTION, FIDELIS XPS COMMANDPOST, the FIDELIS SECURITY SYSTEMS logo and/or other FIDELIS SECURITY SYSTEMS products referenced herein are trademarks of Fidelis Security Systems, Inc. All other trademarks and copyrights contained in this document are owned by their respective trademark & copyright holders. Copying, reuse or distribution of any material contained herein is expressly prohibited. Copyright © 2010 Fidelis Security Systems, Inc. All rights reserved.

Introduction

The information technology landscape is a battlefield. Ready or not, you've been drafted into an epic good vs. evil struggle. You're one of the good guys, striving to protect your networks and data from misuse. Meanwhile, the bad guys poke and prod at all of your assets from within and without, looking for the most profitable ways to exploit the holes in your defenses. They do this patiently, mercilessly, inexorably—whether they're nation-states advancing a political or economic agenda or transnational criminal gangs simply in it for the money. In this context we will take a brief look at today's threats, and then we'll examine the countermeasures enterprise networks have deployed in response to those threats. We'll assess the current effectiveness of the usual mix of defenses and finally we'll introduce you to Fidelis XPS™ core architecture, Deep Session Inspection™, a targeted and uniquely effective technology weapon that's ready to bring your information security arsenal up to the challenges brought by today's clever and determined adversaries.

Advanced Persistent Threat. It's a new term for an old problem.

The Threats

"Advanced Persistent Threat." It's a new term for an old problem. The tactics and techniques for penetrating your defenses aren't much different from what hackers have done since the days of Morris and Michelangelo¹. Online fraudsters and the hackers and script kiddies who preceded them haven't substantially changed their strategy: they still identify human and non-human targets based on a variety of factors; they still probe, surveil, and reverse-engineer their targets to see how they operate and how they can be tricked or defeated; and they still abuse their targets in fulfillment of their goals and objectives.

Consider a few statistics and headlines:

- "Cost of defending against hackers estimated at \$10 billion annually"
- "Last year the Pentagon was subjected to approximately 250,000 online attacks"
- "Credit card fraud: one breach resulted in over \$10 million in losses"
- "70% of breaches originate from inside the organization"

¹ I refer to the 1992 virus, not the polymath. Were there hackers in the sixteenth century? I look forward to your e-mails.

Sound familiar? These factoids and statistics were ripped from the headlines—in 1998.²

What's new clearly isn't the evolved strategy of the bad guys so much as it is a continuation of the trend established in the mid 1990's. The stakes continue to rise, and are higher than ever in 2010: the world's most prominent corporations and governments have made the Internet completely crucial to the way they do business. Everything from operations to finance is happening online. The same is true of Aunt Sally: more than ever she's paying her bills, staying in touch, and managing her retirement portfolio over the Internet. Yet we're still running vulnerable applications and operating systems, still coping with employees we cannot completely trust (their misdeeds may or may not arise from malice), still counting on the same security vendors³ to protect us from adversaries who will never, ever stop innovating.

So the bad guys advance upon us, more steadfastly and profitably than ever. But because of our increased reliance on the Internet for everything we do, the financial opportunities for criminal gangs, cyberterrorists, and state-sanctioned organizations to do us harm are greater than ever; gone are the days when hackers defaced websites merely for bragging rights, or compromised weak computer systems just because they could. Now our adversaries break our applications, they hack into our operating systems and they trick our employees (and our Aunt Sally), for stakes that are much, much higher. Most importantly, they're using techniques that make them nearly invisible to today's most widely used countermeasures.

It's not enough to deliver application awareness when the content itself is the threat.

The Defenses

In response to early Internet threats, an assortment of network-based countermeasures arrived: Firewalls, Intrusion Detection & Prevention systems (IDPS), antivirus, web filtering, and—most recently—Data Leakage Prevention (DLP). Are they measuring up to the challenge? Can an enterprise deploy these tools and rest easy?⁴

Firewall

The first network firewalls arrived in the late 1980's and are, of course, ubiquitously deployed. Firewall technology has improved in the 22 years since then, but little has changed since firewall wizards Marcus Ranum, Matt Curtin and Paul D. Robertson

² <http://attrition.org/errata/statistics/archive.html>

³ See also Joshua Corman's blog post, "Do the Evolution...", posted last year and as relevant as ever: <http://fudsec.com/do-the-evolution-1>

⁴ Spoiler alert: no.

warned against overestimating their impact on business risk in their “Firewall FAQ,” first posted in 1995:

Firewalls can't protect against tunneling over most application protocols to trojaned or poorly written clients. There are no magic bullets and a firewall is not an excuse to not implement software controls on internal networks or ignore host security on servers. Tunneling “bad” things over HTTP, SMTP, and other protocols is quite simple and trivially demonstrated. Security isn't “fire and forget.”

Lastly, firewalls can't protect against bad things being allowed through them. For instance, many Trojan Horses use the Internet Relay Chat (IRC) protocol to allow an attacker to control a compromised internal host from a public IRC server. If you allow any internal system to connect to any external system, then your firewall will provide no protection from this vector of attack.⁵

Read that last sentence again. *If you allow any internal system to connect to any external system*, information will be exchanged that is invisible to every firewall, and unless that information is examined for threats a critical attack surface will remain unguarded. Even so-called “next-generation” firewalls—while they deliver useful improvements over third-generation stateful firewalls, such as identity awareness and application recognition—fall short when it comes to payload and document decoding. If the bad guys are sending your employees zipped renamed PDF documents that contain polymorphic, obfuscated malcode⁶, then it's going to take more than a firewall to respond to those threats (which now arrive with the content, invisible to packet inspection technologies). It's not enough to deliver application awareness when the content itself is the threat.

what innovations has your packet-based IDPS vendor delivered in response to voluntarily downloaded, deeply embedded content as an infection vector?

Intrusion Detection & Prevention

There was a time not long ago when information security success or failure seemed to hinge solely upon our ability to protect our servers from never-ending intrusion attempts. Fortunes were made on the promise that one might someday be able to correlate known applications (or better yet, confirmed host-specific vulnerabilities) on an enterprise's network with incoming attacks targeting those systems, thereby thwarting the attacks and defeating the bad guys. Joanne Cummings considers this exciting possibility in *Network World*:

⁵ <http://www.interhack.net/pubs/fwfaq/>

⁶ They are. See “IBM X-Force® 2010 Mid-Year Trend and Risk Report.” <http://goo.gl/k2Z8>

Sounds good. But before such a scenario can occur, two big problems need ironing out. Intrusion-prevention vendors have to find a way to eliminate false positives, and they have to figure out how to run the devices inline without creating network bottlenecks.⁷

Ms. Cummings' article ran in September 2002. How well has your intrusion prevention provider met this challenge in the eight years since? More importantly, what innovations has your packet-based IDPS vendor delivered in response to the now-common use of voluntarily downloaded, deeply embedded content as an infection vector?

Antivirus

Much has changed since the first signature-based antivirus solutions were offered, in the same decade when network firewalls made their first appearance. Hundreds of millions of endpoints have antivirus software installed, yet infections continue at prodigious rates. The reasons for this are legion. Modern viruses are virtually immune to signature-based approaches, because they have no fingerprints: they change forms as they propagate. Frequently viruses (and their variants) spread faster than signature updates can be deployed to stop them; it's notoriously difficult to propagate new pattern files quickly enough to make a difference⁸. And the sheer volume of unique virus samples that must be analyzed by AV researchers continues to grow exponentially. Signature-based virus defenses have not scaled with the threat and are not providing us with the protection we need. Behavior-based technologies have shown some promise but face serious challenges: can they prevent infections from arriving over the network? Can they do it at multi-gigabit speeds? Can they provide protection against embedded threats on arbitrary ports and protocols? Can they spot customized, targeted attacks? Can they deliver protection without excessive tuning or false positives? Antivirus has never been—and never will be—a security silver bullet.

Web Filtering

Several products aim to protect users while they browse the web. Some perform double duty, protecting against threats while ensuring that employees only visit websites deemed acceptable or productive. Store-and-forward technologies like web proxies can be well suited to detailed inspection and prevention, given several important caveats. For example, inspection latency must be tolerable to users, detection algorithms and policies must be well matched against the threats, and access controls must be in place

⁷ The article is a nostalgic read: <http://www.networkworld.com/buzz/2002/intruder.html>

⁸ "Cyveillance testing shows that even the most popular AV signature-based solutions detect on average less than 19% of malware threats. That detection rate increases only to 61.7% after 30 days."
<http://goo.gl/UGKB>

to ensure that no one has the ability to access the Internet “around” the proxy, circumventing security controls entirely.

Of course, even if all of these criteria are satisfied, web security solutions don’t provide adequate coverage for content tunneled over HTTP, nor can they help manage the flow of information traveling on the 65,533 *other* ports available for applications—rogue or otherwise—to exchange data over the Internet.

Data Leakage Prevention

Data leakage prevention (DLP) has received considerable attention as a separate product category since the term became popular a few years ago. But as information security threats increasingly focus on content, it’s not sensible to approach DLP and information security as satisfying discrete requirements. In the 1995 FAQ quoted above, Ranum probably thought he was stating the obvious when he observed that “many corporations that connect to the Internet are very concerned about proprietary **data leaking** out of the company through that route [*emphasis added*].”⁹ Clearly the first firewalls were designed and deployed with DLP in mind.

It’s not sensible to approach DLP and information security as satisfying discrete requirements. Clearly the first firewalls were designed and deployed with DLP in mind.

A few years after the emergence of DLP as a product category, and two decades after it was understood that the risk of data leakage comes with doing business on the Internet, the DLP market is filled with products that fall short. Some content monitoring & filtering solutions may have excellent document analysis capabilities, but cannot provide real-time prevention on high-speed networks because of their reliance on third party decoding technologies. Other products may deliver excellent results protecting popular channels like web, mail or instant messaging (IM), but suffer from architectural limitations preventing them from scaling to a broader set of port-independent protocols at gigabit speeds. While traditional DLP tools may be helpful for demonstrating regulatory compliance or identifying broken business processes, their performance limitations and focus on well-known communication channels make it hard for them to truly protect intellectual property, enable secure outsourcing, or intelligently manage employee use of social networking services.

⁹ <http://www.interhack.net/pubs/fwfaq/>

It's Not Working

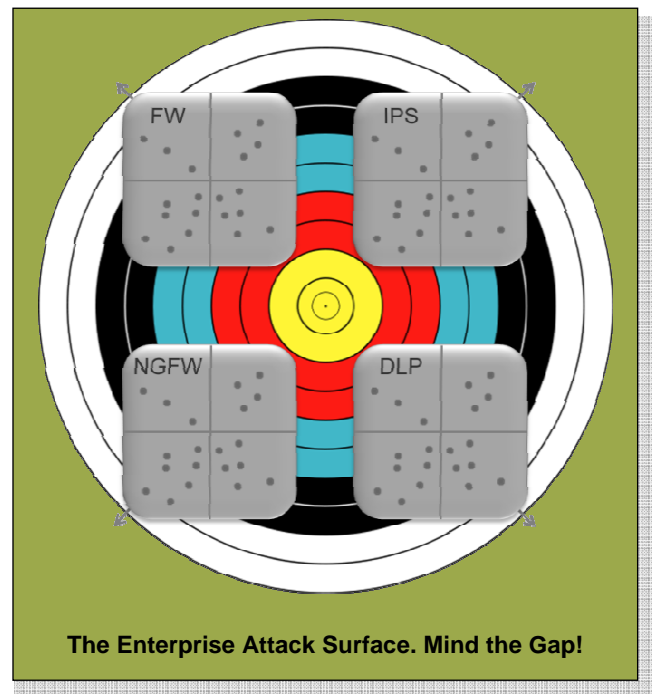
It would certainly not be sensible to simply throw out all of these technologies and start over. Each of them provides a measure of protection that is better than nothing. But neither can we ignore these limitations in light of today's advanced, evolved threats: the bad guys know at least as much as we do, and are adept at sidestepping the outdated defenses offered by today's packet-based countermeasures.

Legacy Technology	Vulnerable To
Firewall	Content-based threats
Intrusion Prevention	Deeply embedded malware, client-side attacks
Antivirus	Highly polymorphic, custom and semi-custom threats
Web Security Proxy	Port- and protocol-agile threats
First Generation DLP	Extrusion/extraction of sensitive information on non-standard ports

Another barrier to your success has to do with market forces. Consider the progression of available products, all marketed to help customers manage IT risk. So many technologies! Firewalls lacked the granularity to defend against certain kinds of intrusions, and so the intrusion detection/prevention market was born in the nineties. As we entered the aughts¹⁰, existing security controls didn't seem up to the task of defending personally identifiable information (PII) or corporate secrets against unauthorized disclosure and so the DLP market emerged. Traditional firewalls lacked application awareness, creating an opportunity for startups to define and occupy a "next generation" firewall market. Signature-based antivirus vendors continue to face the monumental task of analyzing and cataloging a never-ending stream of malware samples, leaving additional opportunities for innovators on both sides of the information security struggle. As you scramble to cover your attack surface, two things seem to be happening simultaneously: security technologies become less effective over time, and security technology vendors cluster their offerings into easily defined product categories. As a result, innovations within these categories are mainly driven by competition among products in each space, overlooking wider trends in the threat landscape and leaving gaps for the bad guys to find and sneak through.

¹⁰ What do *you* call a century's first decade? I look forward to your e-mails.
<http://en.wiktionary.org/wiki/aughts>

For an even more vivid illustration of the weaknesses inherent to today's network countermeasures, look no further than the growing prominence of virtualization and cloud computing. As enterprises make and implement plans to reduce costs through virtualization and application outsourcing, traditional network defenses grow ever more powerless to monitor and regulate the flow of sensitive information and content-borne malware. But to survive in this changing landscape you'll need to do just that. "Measure what is measurable, and make measurable what is not so," said Galileo. Customers today tend to put this inclination more bluntly: "we don't know what we don't know," they tell us, and this challenge will intensify as the cloud trend continues. Fidelis XPS *makes it measurable*, even in the cloud.



Mind the Gap: Fidelis XPS' Deep Session Inspection

In the context of the challenges above we present Fidelis XPS and its patented¹¹ Deep Session Inspection architecture. Fidelis XPS provides visibility and control where other products cannot, with real-time inspection and prevention all the way from low-level network flow metadata up through protocols, applications, and most importantly—*content*. This awareness of content and context—together with a powerful and user-friendly policy engine—provides enterprises with the power to prevent sophisticated breaches and malware infestations. We'll go into more policy detail later, but here's a quick example: with Fidelis XPS you can mitigate the risks of social networking applications and enforce a policy banning Facebook, except for the marketing department, who is only permitted to use the official corporate Facebook account during business hours (but can never chat, send private messages or access Facebook apps hosted overseas). This level of policy granularity, combined with our real-time prevention-enabled performance, is what puts Fidelis XPS in a class by itself.

Other network security tools either lack real-time prevention-enabled awareness, policy-addressable visibility into deeply embedded content, or both.

Why Session Analysis Matters



Network communication is all about encapsulation. Session layers are like matryoshka dolls: each one must be opened to see the next successive layer. Ethernet frames carry IP packets, which contain TCP sessions, which contain HTTP messages, which can carry a set of compressed MIME payloads, one of which might deliver a ZIP archive containing the document now before you. Never has this encapsulation chain been more relevant to information security than it is today. Whereas Fidelis XPS provides full real-time inspection and prevention at all of these levels—and more, as we shall see—

¹¹ <http://www.google.com/patents?vid=USPAT7467202>

simultaneously, other network security tools either lack real-time prevention-enabled awareness, policy-addressable visibility into deeply embedded content, or both.

Additionally, today's networks carry a diverse set of protocols, which deliver content on an unpredictable set of ports. For example:

- One FTP transfer can consist of several simultaneous network sessions, some of which may transit over randomly assigned ports.
- One Skype call or login event can consist of dozens of distinct TCP or UDP sessions to multiple destinations, even over HTTP or HTTPS ports.¹²
- Facebook, like many other modern web destinations, delivers numerous applications and end-user interactions via multiple simultaneous HTTP sessions, which must be regarded as a single "session" in order to make sense of them.

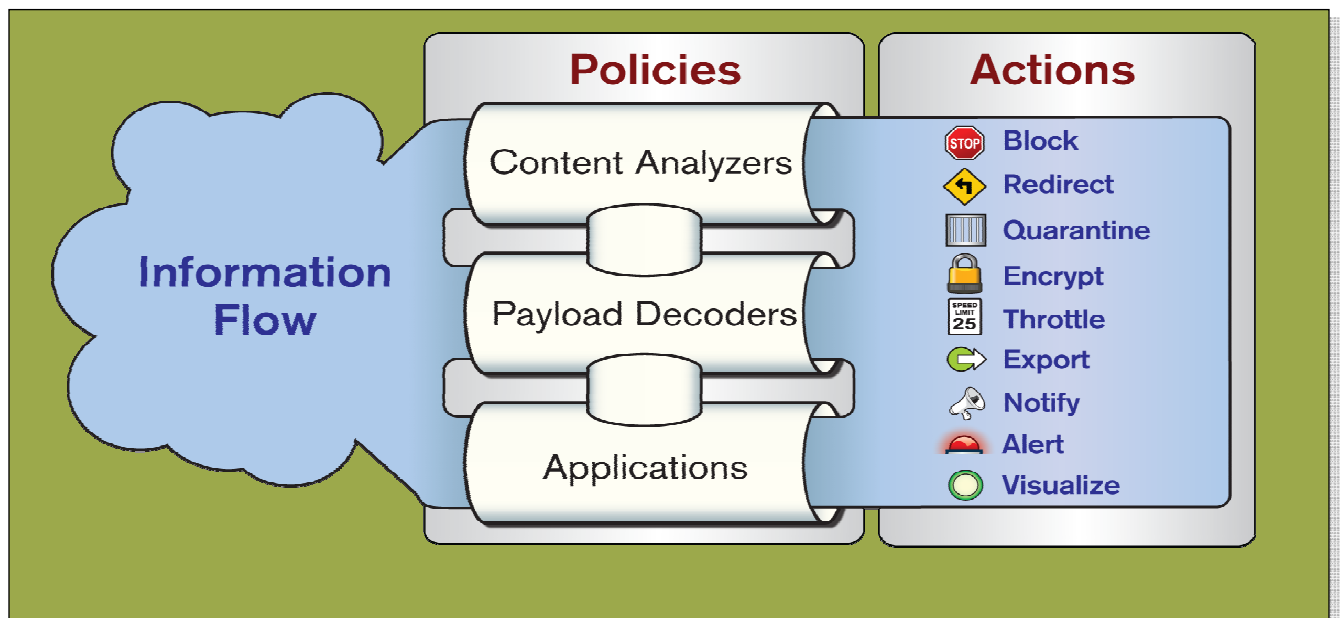
And speaking of the HTTP protocol, it's increasingly being used to tunnel and transmit all kinds of content beyond simple web pages. As one of our customers has observed, "HTTP is now the Wild West."

The challenges don't end with the innermost encapsulation layers: many internal file formats do not clearly represent the information seen or manipulated by an end user. Applications such as Excel or Acrobat Reader may go through all sorts of processes for decoding and presenting files—as a result of information strewn across disparate packets or even sessions—in order to sort out what is displayed on the screen. Without detailed awareness of these obfuscation-enabling, application-level processes, packet-based network monitoring tools are virtually blind to modern threats.

These complexities render mere deep *packet* inspection inadequate to the task of enforcing network policy in real time. In order to achieve visibility and control over modern network activity down to the presentation layer, Deep *Session* Inspection is required, and it's only available from Fidelis XPS.

¹² No, we haven't cracked Skype's encryption scheme. But we can alert and prevent on Skype activity, and all those distinct Skype TCP sessions belonging to a user's Skype login produce a single consolidated alert in our console.

How it Works



The Deep Session Inspection architecture consists of three logical components, working in parallel in Fidelis XPS products: a data collection module, a decoding stack, and a policy engine. The data collection module can receive information from a variety of sources, but the vast majority of Fidelis sensors deployed worldwide are configured to process raw IP packets, taken directly from the wire using a network tap or by sitting in-line at network egress points. Once traffic has been collected, it is then processed by the decoder stack, which is responsible for making sense of the flood of incoming information. Several processes take place in parallel during this step: packets and sessions are correlated, network protocols are identified (regardless of port), application protocols are examined and parsed, and content payloads (if any are present) are decoded and analyzed. Finally, the knowledge extracted from all of the incoming network traffic is subjected to policy enforcement, and network sessions that violate policy can be met with a variety of real-time responses, including prevention.

There are three significant aspects to this process that differentiate Fidelis XPS; they concern visibility, performance, and policy.

1. First, the Deep Session Inspection architecture permits Fidelis XPS customers to regulate network activity on the basis of any combination of details about that

- activity. This includes low-level details about source and destination IP address, session-level details such as time, port number or protocol name, application-level details such as domain or protocol usernames, payload details such as file type or encryption method and finally content details: namely the presence of sensitive, confidential, personally identifiable, or malicious information—or any conceivable combination thereof. No other technology can deliver such extraordinary visibility.
2. Also significant is the fact that a single Fidelis XPS sensor can perform all of these activities in real time, at wire speed, on network segments with throughput up to 2.5 gigabits per second¹³, quickly enough to take an enforcement action by terminating sessions that violate policy. This unprecedented performance characteristic isn't simply due to the fact that our engineers are the best and the brightest; there's a key architectural difference between Fidelis XPS and packet-based products. Deep packet inspection technology was originally created to supplement the capabilities of existing devices like routers and firewalls—devices whose primary function was to receive and retransmit traffic at high speed. These packet processors are optimized for fast analysis of packet headers, all with the goal of routing each packet to its next destination as rapidly and reliably as possible. They were never designed for deep content awareness at wire speed. Nearly ten years ago, Fidelis engineers architected a solution for analyzing embedded content and approached the problem from the ground up, utilizing technology focused directly on content. Session—not packet—analysis is required in 2010 and beyond, and that's where we shine.
 3. The third important advantage provided by Deep Session Inspection has to do with policy, to which we will now turn our attention.

The Power of Policy: Think Fidelis!

Meet Enzo¹⁴, a happy Fidelis XPS customer. Enzo told us recently that in the process of getting the most out of his investment in Fidelis XPS solutions, a turning point came when he started to “think Fidelis.” Fidelis XPS takes an approach to policy that is both unusual and intuitive, but to get maximum value from it one has to “think Fidelis.” Here's how the Fidelis XPS policy engine lets you harness the performance, visibility, and control delivered by the Deep Session Inspection architecture.

¹³ That's 2.5Gbps of guaranteed full-inspection and prevention-enabled throughput per individual sensor, with our tests regularly doubling that number. Need more? Talk to us about our carrier-class solutions.

¹⁴ Not his real name. Security dudes defending high-profile assets tend to be secretive.

Fidelis XPS policy is constructed from building blocks we call fingerprints. These fingerprints come in “who,” “what”, and “how” categories, focused on location (the origin and destination of each session, by IP address, geography, directory attribute or reputation), content (each message itself, after all the protocols, payloads, applications and compression have been stripped away), and channel (any of thousands of attributes having to do with how the message was transmitted). These fingerprints are assembled into rules as Boolean expressions. A rule expression can be simple or complex. If you can describe a network event in words, you can write a rule for it, with assistance from our thoroughly documented and customizable pre-built policies, from our support team, from our other customers, and from the various training options offered by Fidelis.

Dozens of different types of location (“who”), content (“what”), and channel (“how”) fingerprints are available, enabling the creation of a content protection and network security policy that’s as simple or as complex as conditions dictate.

Here’s a trivial favorite:

`Exe AND ImageFileName`

In English: “prevent or alert on any traffic containing an executable that’s been renamed to look like an image.” Remember, this expression doesn’t care if the renamed executable is buried in a dozen nested zip archives; Fidelis XPS will still spot this payload quickly enough to block the transfer, if prevention is specified as part of the rule.

Remember the rather complex Facebook example we gave earlier?

In English:

Facebook is banned, except for the marketing department, who is only permitted to use the official corporate Facebook account during business hours (but can never chat, send private messages or access Facebook apps hosted overseas).

Think Fidelis! If the conditions in a policy expression are met (i.e., if an expression evaluates as “true” at any point during any network session), the rule containing the expression is activated. What conditions on the network must be met in order to violate the Facebook ban and produce an alert or blocking response? There are several valid approaches to constructing a Boolean expression describing these conditions. Here’s one:

`FB AND (NOT (FromMktDept AND CorpFBacct)
OR FBChat OR FBPM OR DstOverSeas OR OutSideBusHours)`

The expression might look daunting, but it's essentially a big "and" statement, with arguments on the left and right of the "and" operator, both of which must be true for the expression's rule to fire. On the left: "FB," a channel (i.e., "how") fingerprint that will evaluate as "true" for any session where an end user accesses the Facebook service. On the right is a set of four criteria joined with the "or" operator. Any of these four conditions would individually constitute a violation of the Facebook rule if they belong to Facebook traffic: 1) either the access isn't coming from the marketing department (FromMktDept), or the Facebook account being used isn't the official corporate account (CorpFBacct); or 2) Facebook is being used for interactive chat regardless of who's doing it (FBchat); or 3) Facebook is being used by anybody for private messaging (FBPM); or 4) the IP address of the destination server is overseas (DstOverseas); or 5) the Facebook activity is taking place outside normal business hours (OutsideBusHours). Each of the fingerprints used to build this expression can be defined and tweaked in the Fidelis XPS policy editor, providing an extraordinary degree of visibility and control; remember that these rules can be used to implement real-time detection and prevention.

Dozens of different types of location ("who"), content ("what") and channel ("how") fingerprints are available, enabling the creation of a content protection and network security policy that's as simple or as complex as conditions dictate. These "who," "what" and "how" fingerprints provide access to all of the decoding algorithms in the Fidelis XPS Deep Session Inspection architecture, providing exceptional levels of granularity as you select and refine the policies most suited to the needs of your enterprise¹⁵.

Extraordinary Possibilities

Now that Enzo "thinks Fidelis," he has used the Deep Session Inspection architecture to great effect, isolating and eradicating over 400 botnet infections on his network using efficient techniques that are impossible with any other product. Indeed, the infections he's detected and fixed using Fidelis XPS this year went completely undetected by the firewall, intrusion prevention and antivirus technologies pervading his network. Here are some of the tricks Enzo used—and continues to use—to keep an eye out for bad behavior on his network:

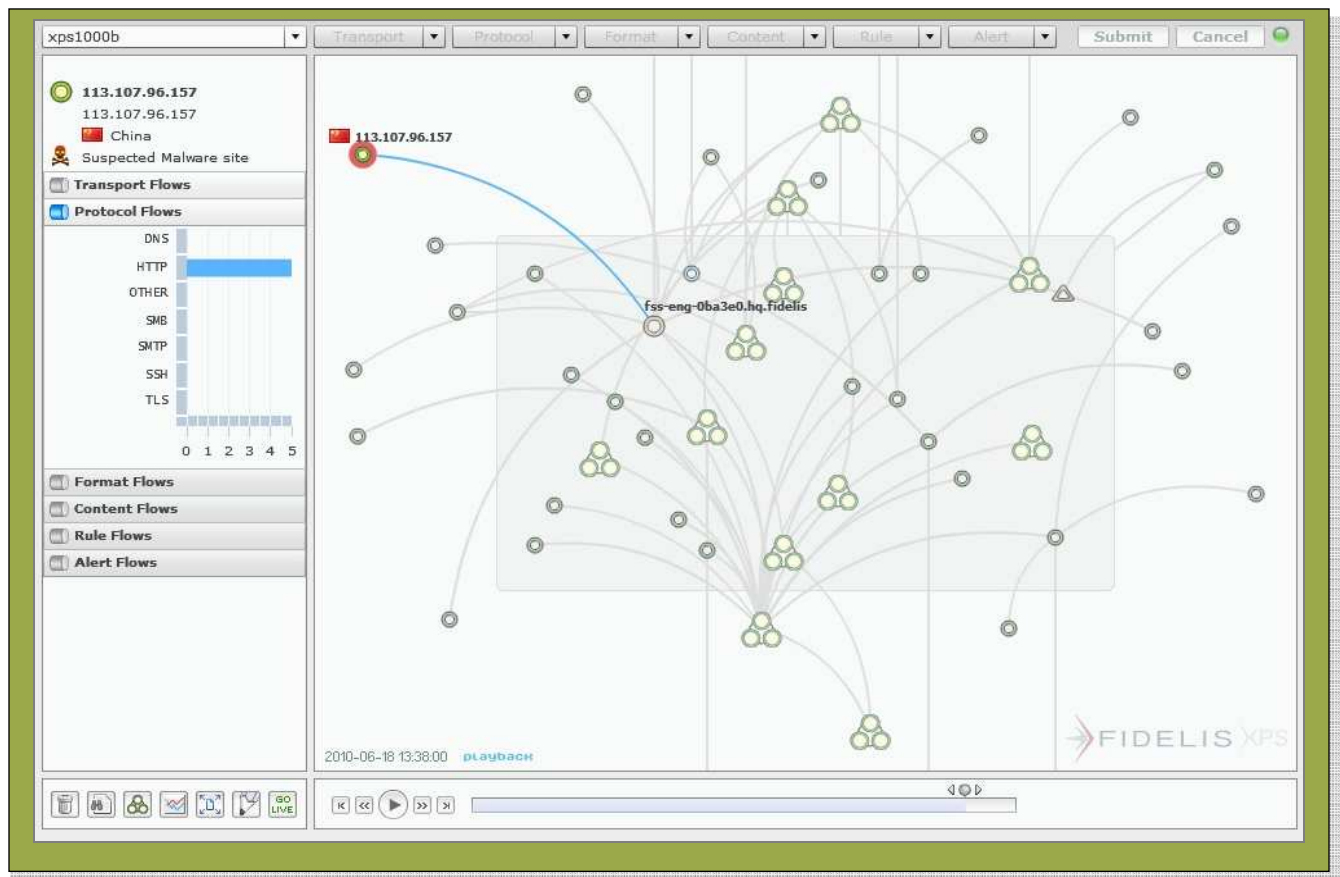
- Traffic on port 80 or 443 that doesn't look like HTTP;
- Traffic destined for suspicious countries, particularly encrypted traffic or sessions taking place on weekends or outside business hours;
- Short IRC messages (these can often be botnet command & control activity);

¹⁵ Six location algorithms. Eleven channel algorithms. Thirteen content recognition algorithms. Thirty seven document type families. Fifty eight network and application protocol types. The list is always growing. We love talking about this stuff. Drop us a line and we can look at our policy elements with you in more detail.

analysis. To download and examine the PDF document near the bottom of the decoding process, click it. To download and analyze the zip file one layer up, click it! Each alert triggers a capture of full network session data, opening a range of possibilities we call *intelligent network forensics*.

Information Flow Map: See Everything

What about all of the network traffic that doesn't break any rules or violate any known policies? Recall that every network session is subjected to detailed analysis by the Deep Session Inspection engine. Every protocol is analyzed, each payload decoded and examined. A wealth of knowledge is available as a result of this process, even—especially—when it comes to network flows that don't break specifically defined policies. The Fidelis XPS Information Flow Map™ presents information about these flows in real time.



This interactive map, a standard feature in Fidelis XPS sensors, provides the whole picture from source and destination metadata up to payload details, filterable and adjustable to show specific protocols, payloads or content. Everything the Deep Session Inspection engine uncovers as it examines network traffic is available in this view, enabling customers to prioritize their network policy development, enforcement and incident response activities. For more information about Fidelis XPS' Information Flow Map feature and how you can use it to see your network in a new light, check out some of our videos on the subject.¹⁶

Operationalizing Dynamic Threat Intelligence

Fidelis XPS network security appliances also include the Feed Manager feature, which arrives preconfigured with an expanding array of threat intelligence feeds. These dynamically updated sources of reputational knowledge currently focus on malware and phishing prevention, and on recognition of traffic destined for known botnet command and control (C&C) networks. These advanced data sources, continuously updated by our threat intelligence partners and automatically kept up to date by Fidelis XPS, allow our customers to enforce network security policy by monitoring and controlling network access to devices known to host phishing, malware or C&C threats. Combined with the Deep Session Inspection engine and the other policy controls available in Fidelis XPS, these feeds provide risk mitigation that's independent of the malware obfuscation innovations that keep security managers up at night. Fidelis is closing the gaps left open by legacy network security products.

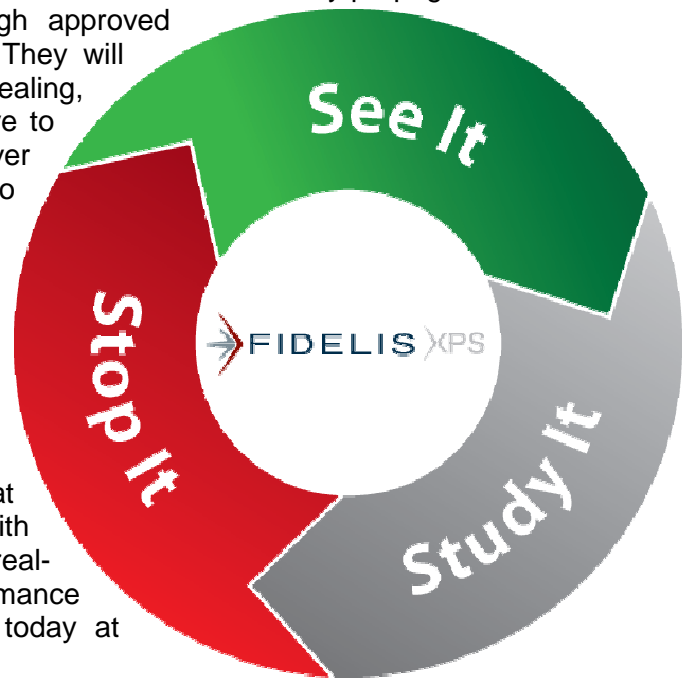
Fidelis is closing the gaps left open by legacy network security products.

Conclusion

By now it should be clear that Fidelis XPS is not another packet-focused network security solution such as a firewall, IPS, AV, or web filtering product. And it's definitely more than just another DLP product. It's a powerful network security solution that delivers real-time visibility and control over all of the traffic crossing your network. Our patented Deep Session Inspection architecture, the Information Flow Map feature, a powerful and novel policy engine, unsurpassed reporting capabilities, and the threat intelligence feeds available from Fidelis all combine seamlessly to help you fill in the gaps left by the other tools in your arsenal—giving you the situational awareness you need to actually do something *proactive* about today's advanced threats.

¹⁶ <http://youtube.com/FidSecSys>

The bad guys have infection vectors that are immune to IDPS. They propagate malware that's impervious to AV, tunnels through approved channels and is resistant to analysis.¹⁷ They will keep plotting, innovating, probing and stealing, and new buzzwords will continue to arrive to describe their evolved approach. Whatever "it" may be, the same old tools are no longer enough. With Fidelis XPS you'll have a better chance to See it, Study it, and Stop it. You'll **see** everything on your network through our Information Flow Map feature. You'll **study** the threats and easily distinguish between good traffic and bad traffic through intelligent network forensics and our novel policy engine, augmented by our up-to-the-minute threat intelligence feeds. And you'll **stop** it with confidence thanks to our unprecedented real-time prevention-enabled performance characteristics. To learn more, call us today at 1.800.652.4020.



¹⁷ "Emerging Qakbot Exploit Is Ruffling Some Feathers," Dark Reading, 10/26/2010. <http://goo.gl/J62K>