
From: "Penny Leavy-Hoglund" <penny@hbgary.com>
To: "Ted Vera" <ted@hbgary.com>; "Aaron Barr" <aaron@hbgary.com>
Sent: Thursday, April 15, 2010 6:36 AM
Attach: 451_ESP_E-Crime_and_APTs_Final.R1.pdf
Subject: FW: 451Group eCrime Report Published

[Here it is, GREAT report](#)

From: Karen Burke [mailto:karenmaryburke@yahoo.com]
Sent: Friday, April 09, 2010 8:11 AM
To: penny@hbgary.com
Subject: 451Group eCrime Report Published

Hi Penny, I just got a copy of The 451Group eCrimes report. I haven't read it yet -- I wanted to pass on to you first. I'll read it today. Rajeev is providing free access --below is a note from Rajeev to me since I had trouble accessing it using our client login.

Hi Karen, glad to see the onsite is coming up in the near future.

Regarding the report, the reason you weren't able to access it is that these reports price out as \$3750 purchased ad-hoc - if they're not built into the agreement in advance. Given HBGary's involvement in this one however, I'm keen on just giving you the PDF with the intention of this helping our efforts to renew the relationship once the 6-month term is up.

Report attached and you're now forewarned that I'll be bringing this up come renewal time ;)

cheers,

Rajeev Pattni
Senior Account Executive
617-261-1660 – (o)
617-699-2952 – (m)
[The 451 Group](#) | [Tier1 Research](#)
Analyzing the business of enterprise IT innovation

E-CRIME AND ADVANCED PERSISTENT THREATS

How Profit and Politics Affect IT Security Strategies

Cybercrime and sophisticated state-sponsored hacking are forcing enterprises to search for new approaches to securing their networks and endpoints that frees them from the 'whack a mole' game they're stuck in. What's needed are tools for spotting sophisticated crime patterns and thwarting them.

ESP | ENTERPRISE SECURITY PRACTICE

4 FINDINGS

- Sophisticated cybercriminal attacks focused on saleable data and state-sponsored hacks aimed at state secrets or valuable IP are increasingly the focus of IT security efforts at firms in verticals like government, energy, finance and technology. **PAGE 5**
- Existing IT security investments such as endpoint anti-malware, IDS/IPS and firewalls are necessary but insufficient to detect and block modern threats and protect enterprise data. **PAGE 8**
- The advent of advanced persistent threats like those targeting defense contractors and high-profile IT firms drives demand for capabilities such as threat correlation, reputation monitoring and forensics. **PAGE 12**
- The ability to respond to new threats and attacks is hampered by a lack of reliable, impartial data, and by regulatory compliance which has supplanted security as a main driver of IT security investment. **PAGE 22**

5 IMPLICATIONS

- Organizations must realign their investment in IT security to address the threat posed by professional cybercrime groups and APTs. **PAGE 35**
- Stronger legal frameworks need to be established to enable the investigation and prosecution of cybercrime cases across borders. **PAGE 32**
- Greater investment in areas such as rights management, reputation monitoring, fraud detection, threat correlation and cyberforensics can provide insight into new attacks. **PAGE 52**
- Enterprises need to reevaluate the approach to endpoint protection as the protection offered by multifunction anti-malware suites declines. **PAGE 37**
- Improved network monitoring, analysis and incident response are needed to battle sophisticated threats and data theft. **PAGE 47**

1 BOTTOM LINE

Organized cybercrime and APTs aren't new, but they now pose a much bigger threat to the safe conduct of commerce and to public safety and national security. Sadly, most enterprises are still fighting the last war against loud, dumb attacks like Code Red, Blaster and Slammer. The gulf between protection and threat is wider than ever. Enterprises need to redirect their security investment to products and services that address the new IT security reality: increasing their ability to capture, analyze and understand network flows, to monitor threats specific to their company or vertical, and to distribute and enforce granular use and access policies that limit risk both inside and outside the network firewall. There are no easy fixes, and IT vendors cloud the discussion with FUD and recommendations tailored to their product lines.

MARCH 2010

ABOUT THE 451 GROUP

The 451 Group is a technology analyst company. We publish market analysis focused on innovation in enterprise IT, and support our clients through a range of syndicated research and advisory services. Clients of the company — at vendor, investor, service-provider and end-user organizations — rely on 451 insights to do business better.

ABOUT TIER1 RESEARCH

Tier1 Research covers consumer, enterprise and carrier IT services, particularly hosting, colocation, content delivery, Internet services, software-as-a-service and enterprise services. Tier1's focus is on the movement of services to the Internet — what they are, how they are delivered and where they are going.

© 2010 The 451 Group, Tier1 Research and/or its Affiliates. All Rights Reserved. Reproduction and distribution of this publication, in whole or in part, in any form without prior written permission is forbidden. The terms of use regarding distribution, both internally and externally, shall be governed by the terms laid out in your Service Agreement with The 451 Group, Tier1 Research and/or its Affiliates. The information contained herein has been obtained from sources believed to be reliable. The 451 Group and Tier1 Research disclaim all warranties as to the accuracy, completeness or adequacy of such information. Although The 451 Group and Tier1 Research may discuss legal issues related to the information technology business, The 451 Group and Tier1 Research do not provide legal advice or services and their research should not be construed or used as such. The 451 Group and Tier1 Research shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice.



*Analyzing the Business
of Enterprise IT Innovation*



Better perspective from the top in independent tech research

New York

20 West 37th Street, 6th Floor
New York, NY 10018
Phone: 212.505.3030
Fax: 212.505.2630

San Francisco

140 Geary Street, 9th Floor
San Francisco, CA 94108
Phone: 415.989.1555
Fax: 415.989.1558

London

37-41 Gower Street
London, UK WC1E 6HH
Phone: +44 (0)20.7299.7765
Fax: +44 (0)20.7299.7799

Boston

52 Broad Street, 2nd Floor
Boston, MA 02109
Phone: 617.261.0699
Fax: 617.261.0688

TABLE OF CONTENTS

SECTION 1: EXECUTIVE SUMMARY	1
1.1 KEY FINDINGS	3
1.2 METHODOLOGY	4
 SECTION 2: SHIFTING SANDS: CYBERCRIME AND THE ENTERPRISE	 5
2.1 CYBERCRIME – THE NEW FACE OF ENTERPRISE THREATS	5
2.1.1 <i>The Cybercrime Ecosystem.</i>	6
2.1.2 <i>Enterprise Exposure to Cybercrime and Advanced Persistent Threats</i> . . .	7
FIGURE 1: <i>Enterprise Cybercrime Exposure</i>	7
2.1.3 <i>Enterprise Coverage for Cybercrime and Advanced Persistent Threats</i> . . .	8
FIGURE 2: <i>Enterprise Security Coverage</i>	8
2.2 ANTI-MALWARE AND THE ILLUSION OF PROTECTION	9
2.2.1 <i>Gaming the Signature Game</i>	9
2.3 ADVANCED PERSISTENT THREATS	12
2.3.1 <i>APTs in action: Aurora, GhostNet</i>	12
2.3.2 <i>APTs: Common Characteristics</i>	15
FIGURE 3: <i>Feature Checklist – Commodity Malware vs. APT</i>	18
2.4 CYBERWARFARE AND ATTACKS ON CRITICAL INFRASTRUCTURE	19
FIGURE 4: <i>Critical Infrastructure Attacks</i>	20
 SECTION 3: FUD FACTOR: HYPE, COMPLIANCE AND THE PROBLEM OF INFORMATION ASYMMETRY	 22
3.1 SIZING THE CYBERCRIME AND APT PROBLEM.	23
FIGURE 5: <i>Estimated Losses From TJX Data Breach (\$B)</i>	25
3.2 WHACK-A-MOLE PAYS THE BILLS	26
3.3 FUD AND ITS DISCONTENTS.	27
3.4 THE PROBLEM OF INFORMATION ASYMMETRY	28
FIGURE 6: <i>Information Asymmetry.</i>	29
3.5 A COMPLIANCE CONUNDRUM	30
3.6 FIXING LAWS AND POLICY	32

SECTION 4: AN ENTERPRISE CYBERCRIME TOOLKIT 35

4.1 HARDENING THE ENTERPRISE ENDPOINT	37
4.1.1 PC Lifecycle Management Subsumes Threat Protection	38
4.1.2 Application Control and Beyond.	39
4.1.3 A New Endpoint Paradigm: 'Presumed Owned'.	40
4.2 IMPROVED THREAT INTELLIGENCE AND CORRELATION	42
4.2.1 Broader Applications for Anti-Fraud	45
4.3 LEVERAGING NETWORK MONITORING AND ANALYSIS	47
4.3.1 Growing Eyes and Ears: Network Traffic Analysis, DLP and Beyond	47
4.3.2 Drilling Down on Threats: Improving Incident Response.	49
4.3.3 Pulling It All Together: SIEM and Event Correlation	50
4.4 PROTECTING DATA: A SHIFT TO INFORMATION RIGHTS MANAGEMENT	52

SECTION 5: M&A OPPORTUNITIES 54

5.1 BUILDING A CYBERCRIME-FIGHTING STACK	54
5.2 THREAT INTELLIGENCE FIRMS FIND SUITORS IN THE WINGS	54
5.3 MALWARE FORENSICS: EVERYONE'S SECRET SAUCE	55
5.4 SIEM TAKES CENTER STAGE	56
5.5 OPPORTUNITIES FOR (SMARTLY) MANAGED ENDPOINTS.	57
5.6 BOTS ARE HOT (AGAIN)	59

SECTION 6: VENDOR PROFILES 61

6.1 ARCSIGHT	61
6.2 BRIGHTCLOUD	62
6.3 COMMTOUCH SOFTWARE	63
6.4 HBGARY	64
6.5 MARKMONITOR	65
6.6 MEMENTO SECURITY	66
6.7 NETWITNESS	67
6.8 NICE	68
6.9 PALANTIR TECHNOLOGIES	69

6.10 QINETIQ 70

6.11 RSA (THE SECURITY DIVISION OF EMC) 71

6.12 TEAM CYMRU 72

6.13 TRIUMFANT 73

6.14 TRUSTEER 74

INDEX OF COMPANIES 75

SECTION 1

Executive Summary

Great crimes always contain an element of farce. And so it was with the record-busting five-year crime spree conducted by Albert Gonzalez, aka ‘segvec,’ aka ‘cumbajohny,’ aka ‘souponazi’ – the mastermind behind one of the greatest online heists of all time. Gonzalez, we now know¹, managed to carry out the bulk of his crimes while working under the nose of law enforcement as an informant for the FBI, and under the nose of IT administrators and auditors at some of the US’s largest corporations. Those crimes included serial hacks of TJX, OfficeMax, Dave & Busters, Hannaford Supermarkets and then, of course, credit card processor Heartland Payment Systems, a proverbial mother lode that netted Gonzalez information on 130 million US credit card accounts. The fact that Gonzalez’s retinue included a bevy of unnamed hackers “resided in or near Russia” and a seven-foot-tall body-building programmer working in the bowels of Morgan Stanley only added to the dismal amusement of the affair.

For businesses and consumers, the aftershocks of Gonzalez’s hacks will be felt for years. Just the direct losses tied to data breaches are eye-popping. SEC filings from TJX put the cost of that company’s data breach at more than \$200m in fines, restitution and other breach-related matters. That’s a tenfold jump from the \$20m-25m that ChoicePoint estimated its data breach would cost to shareholders back in 2005. Even larger costs will be borne by companies in the form of stricter regulations and the attendant costs for compliance with them. The PCI Council, for example, is already tightening requirements for internal audit and stiffening penalties for compliance. That group was stung by reports that Heartland had, in fact, passed a PCI audit just weeks before malicious software resident on its payments network began exfiltrating credit card information on millions of consumers.

But what is the larger significance of the Gonzalez attack? Or, a better question might be: is there a larger significance of the Gonzalez attack? Indeed, the fact that many of the headline-grabbing hacks of corporate information in the past five years have been neatly tied back to one individual and a somewhat larger network of associates might be seen as comforting news: by arresting and jailing Gonzalez, law enforcement in the US has cut the head off the hydra and shut down a multimillion-dollar racket in stolen credit card information. Some² have, in fact, reasoned that the huge numbers of compromised accounts in the Gonzalez hacks are ‘outliers’ – data points that are distorting the real trends around data theft, which show decreases, not increases, in data breaches. We’re not so optimistic. In fact, recent attacks like the coordinated hack of top US IT firms like Google, Intel and Adobe suggest just the opposite: that sophisticated, profit-motivated cybercrime and state-sponsored hacking are growing problems of which Albert Gonzalez was not a particularly remarkable example – though he was quite successful in achieving his objectives.

1. “TJX Hacker charged with Heartland, Hannaford breaches,” *Wired.com*, August 17, 2009

2. Michael Dahn on his blog, *Chaordic Mind* (<http://chaordicmind.com/blog/tag/stephen-watt/>)

Enterprises, girded against big, noisy threats cooked up by bored teenagers, are instead contending with an entirely different threat paradigm: quiet, stealthy and sophisticated attacks by determined and well-financed opponents. Barring wholesale changes in the way enterprises manage the security of their networks, users and business partners, breaches like Hannaford and Heartland, as well as Aurora, GhostNet and Titan Rain, will become more common in the future, not less – even though we may not be reading about them in the headlines. In fact, one of the first victims of the cybercrime epidemic may be the IT security industry itself, which has been slow to react to a fast-changing threat landscape, ill-equipped to deal with the output of industrialized malware production and incapable of detecting nuanced application-based attacks and custom malicious code of the kind that Gonzalez and hacker-accomplice Stephen Watt unleashed on their victims' networks.

What does that portend for the computer security industry? Change – and lots of it. Indeed, we believe that a fundamental shift is already under way within the IT security industry, as existing approaches to network and endpoint defense (e.g., network firewalls, intrusion detection and antivirus software) show their age, while willy-nilly user- and business-driven adoption of new Web-based applications and social networks outstrips the ability of IT staff to keep up and punches holes in perimeter defenses.

At the 10,000-foot level, cybercrime and sophisticated state-sponsored hacking are forcing security software firms to devise new approaches to securing enterprise networks and endpoints that will extricate them from the 'whack a mole' game they're stuck in now, improving the ability to spot sophisticated crime patterns and thwart them. Failing that, enterprises increasingly need help with forensics – identifying the source and extent of the crime and its impact on their business. Fraud prevention, reputation monitoring, more sophisticated threat intelligence and greater agility within IT departments will all play an important role in fighting cybercrime. Companies that don't have that kind of intelligence will be forced to acquire it or strike up advantageous partnerships.

At the same time, ISPs, enterprises and ISVs alike will have to tackle big, squishy problems that most have been loath to address, including the security of underlying application code and the culpability of unwitting (or willing) employees and insiders in enabling cybercriminal attacks. The Rugged Software Manifesto³, to which The 451 Group's Joshua Corman contributed in cooperation with other thought leaders, is one example of the government- and industry-sponsored initiatives that will be needed to change the culture of lax development practices that contribute to software insecurity and, indirectly, feed the cybercrime problem.

This report takes the measure of the modern cybercrime epidemic, makes some predictions about the direction that sophisticated cybercrime and state-sponsored espionage will take in the next few years, and evaluates existing vendor responses to the shifting

3. <http://www.ruggedsoftware.org/>

threat landscape. Next, we take a look at some of the technologies and tools that we believe will be increasingly important parts of the enterprise cybercrime-fighting toolkit. Finally, we consider the existing vendor landscape and what kinds of partnerships and M&A opportunities the cybercrime epidemic might create in the years ahead.

1.1 KEY FINDINGS

Enterprises are increasingly concerned about sophisticated cybercrime and advanced persistent threats – low and slow-moving attacks that can remove sensitive customer and financial data or intellectual property from their networks. However, existing IT security infrastructure and investments are poorly aligned with the new threat vectors, leaving a protection gap.

- IT security vendors have done a poor job explaining and addressing new threats, proposing mostly traditional remedies (“update your antivirus signatures!”) for fast-evolving threats and conflating ‘compliance’ with ‘security’ – all to the disadvantage of customers.
- To address the new threat landscape, changes are needed. First and foremost, enterprises need to reevaluate their approach to endpoint protection, which has become a security flashpoint in an increasingly Web-based, de-perimeterized computing environment.
- Additional investment in network monitoring and analysis and incident response, as well as threat intelligence and correlation, are needed to improve the ability of both private- and public-sector organizations to spot sophisticated, low, slow-moving attacks aimed at them, as well as threats posed by malicious insiders.
- Improved data protection and rights management are also needed, as attackers increasingly focus on sensitive data and intellectual property in their attacks. Adoption of information rights management tools, which blend elements of data-leak prevention, data encryption, identity management and policy management, is a step toward protecting networks against APTs and other sophisticated threats.
- Shifts in the IT security landscape present opportunities for M&A and consolidation. We think capabilities like malware forensics, threat intelligence, threat correlation (including SIEM) and signature-less malware blocking are all of increasing value given the advent of sophisticated cybercrime and APTs. Look for consolidation and M&A activity in these areas.

1.2 METHODOLOGY

This report on enterprise security is based on a series of in-depth interviews with a variety of stakeholders in the industry, including IT managers at end-user organizations across multiple sectors, technology vendors, managed service providers, telcos and VCs. This research was supplemented by additional primary research, including attendance at a number of trade shows and industry events.

Reports such as this one represent a holistic perspective on key emerging markets in the enterprise IT space. These markets evolve quickly, though, so The 451 Group offers additional services that provide critical marketplace updates. These updated reports and perspectives are presented on a daily basis via the company's core intelligence service – the 451 Market Insight Service. Perspectives on strategic acquisitions and the liquidity environment for technology companies are updated regularly via the company's forward-looking M&A analysis service – 451 TechDealmaker – which is backed by the industry-leading 451 M&A KnowledgeBase.

Emerging technologies and markets are also covered in additional 451 practices, including our Enterprise Security, Eco-Efficient IT, Commercial Adoption of Open Source (CAOS), Infrastructure Computing for the Enterprise (ICE) and 451 Market Monitor services, as well as CloudScape, an interdisciplinary program from The 451 Group and subsidiary Tier1 Research. All of these 451 services, which are accessible via the Web, provide critical and timely analysis specifically focused on the business of enterprise IT innovation.

This report was written by Paul Roberts, Senior Analyst, Enterprise Security.

Any questions about the methodology should be addressed to Paul Roberts at:

paul.roberts@the451group.com

For more information about The 451 Group, please go to the company's website:

www.the451group.com

SECTION 2

Shifting Sands: Cybercrime and the Enterprise

2.1 CYBERCRIME – THE NEW FACE OF ENTERPRISE THREATS

In its semiannual Security Intelligence Report for the first half of 2009, researchers at Microsoft took the opportunity to reflect on a decade's worth of malware – from 1999 to 2009. After a soft-focus look at computing in the mid-to-late 1990s – a period before Internet access was ubiquitous and when computer viruses and worms, like the 1997 outbreak of the Melissa macro virus, were still a novelty – MSRC researchers took aim at the latest phase in the evolution of threats to computing systems. This period, which they say began around 2004, is characterized most by what Microsoft terms “profit-oriented malware” – worms like MyDoom, Bagle and Nuwar, whose purpose was to construct vast bot networks of compromised hosts that could be leveraged as commercial platforms⁴ for the distribution of spam, denial-of-service (DOS) attacks and other ‘services.’ The company then goes on to note the considerable efforts that Microsoft has made to secure its applications and operating system, and its cooperation with legal and community-based partners in law enforcement and government to stamp out organized crime. Things are getting better all the time, it would seem!

While it's hard to argue with Microsoft's analysis of the evolution of malware, it's harder to be as sanguine about the prospects for squelching out a growing epidemic of what has been termed ‘cybercrime,’ ‘cyberespionage’ and APTs such as state-sponsored hacking. Indeed, our research – as well as the separate work of researchers, journalists and law enforcement around the globe in recent years – points to a trend that's heading in the opposite direction: the e-criminal enterprise is advancing rapidly, expanding its reach and the sophistication of online criminal conspiracies. More than the companies they target, or even the security firms arrayed against them, online crime groups of all sizes are investing in R&D and embracing new technologies and delivery models (including SaaS, managed offerings and outsourcing) as a way to gain a competitive edge.

As part of our research, we've spoken with senior security officers in verticals hard hit by online crime and fraud, as well as researchers, technologists and law enforcement officials who are working to combat the problem. None of them suggested that the battle was nearly won, nor that there were simple solutions to the problem of online crime. Rather, combating organized online criminal groups – just like their predecessors in the physical world – will require time, patience and legal reforms, as well as coordination between affected companies and industries. There's evidence that some of those changes are coming to pass. But it took more than five decades to break the back of the mafia in the United States – a country with a stable economy, uniform legal code and clear jurisdictional lines. Why should it take any less time to break up vast international criminal conspiracies that

4. Microsoft Security Intelligence Report, Vol 7, p.22.

straddle borders and migrate rapidly to take advantage of countries with lax oversight, corrupt officials and laws that are inadequate to prosecute computer crime?

The purpose of this report, then, isn't to break down doors and shine a light into the dens of online crime syndicates – there are others who are far better suited to that task. Rather, it is to look at how we think cybercrime trends (as well as countervailing forces, such as compliance demands) will impact enterprises and engender changes to the way they secure their data, networks and users, as well as the investments they make (or don't make) in technology to help them do so. We see large changes on the horizon, as enterprises come to grips with the gaping holes left by their existing security investments. We're hoping that this report points the way toward smarter investments in the future.

2.1.1 THE CYBERCRIME ECOSYSTEM

The online crime ecosystem is a diverse one. It ranges from individual scammers looking to make a fast buck off of spam-backed marketing or identity theft, to small criminal gangs such as the one headed by TJX and Heartland hacker Albert Gonzalez, to large, well-organized and diversified syndicates like the former Russian Business Network, which rake in untold wealth from a long list of illicit businesses and online scams that operate globally. Servicing them are countless cottage industries that have sprung up in the past decade for everything from malware authoring to the production and sale of phishing toolkits and phony security software (a burgeoning industry). There's money to be made offering bulletproof ('no questions asked') hosting, botnet rentals and money mule management.

While many of these scams are targeted, broadly, at consumers rather than businesses, there is considerable overlap into the domain of enterprise security, and enterprises large and small are increasingly finding themselves the targets of online scams. Let's look at a couple of notable examples.

2.1.2 ENTERPRISE EXPOSURE TO CYBERCRIME AND ADVANCED PERSISTENT THREATS

One useful exercise is to take a look at the landscape of online crime and to connect the dots between those crimes and enterprise exposures.

FIGURE 1: ENTERPRISE CYBERCRIME EXPOSURE

	SOPHISTICATION	VECTOR	ENTERPRISE EXPOSURE
ATTACKS/ DOS	High	Internet, network	Yes – Enterprises are the primary target of attacks for hire. Online criminals typically target e-commerce firms, gaming and other firms whose business depends on high availability. Enterprise systems enrolled in botnets may also be unwitting participants in criminal online attacks.
DATA THEFT	Medium/high	Email, Web, network, analog, social, physical	Yes – a leading exposure for enterprises across verticals, with particular concern in IP-sensitive industries (tech, pharma, defense, finance).
EXTORTION	High	Email, Web, social, physical	Yes – Enterprises increasingly face risk of extortion at the hands of former employees or business partners, often subsequent to data theft (or legitimate data sharing).*
FENCING	High	Email, Web, social, physical	Yes – Exposure to loss is highest in retail and banking, but any company with a product to sell has exposure to bogus orders stemming from online fencing operations.
IDENTITY THEFT	Medium	Email, Web, network, analog, social	Yes – while identity theft scams are typically aimed at individuals, small and midsize businesses are an increasingly popular target. Financial services is particularly hard hit. Generally, attacks against employees may also yield credentials and sensitive data belonging to their employer, creating additional exposure.
PHISHING	Low	Email, Web	Yes – as with spam, phish attacks are a vector for malware infections within the enterprise. In addition, firms targeted in phish attacks face both monetary and reputational risk. Spear phishing and ‘whaling’ attacks against high-level or critical employees are often the opening gambit used by APTs.
PORN/ ILLCIT	Low	Email, Web	Yes – Illicit websites can be a vector for malware. Employee access to illegal pornographic material or illicit materials can pose reputational and operational risks to their employer as well.
SPAM	Low	Email	Yes – spam runs target indiscriminately; spam impacts availability and productivity and is a vector for malware.

* See the University of Cambridge’s “A Pact with the Devil” Technical Report for a good example of this.
<http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-666.html>

2.1.3 ENTERPRISE COVERAGE FOR CYBERCRIME AND ADVANCED PERSISTENT THREATS

If enterprises are increasingly the targets of sophisticated cybercrime, a natural follow-on question is ‘how well are they protected by their existing investments?’ Here’s a quick and dirty assessment.

FIGURE 2: ENTERPRISE SECURITY COVERAGE

	PERCEIVED PAIN	ACTUAL PAIN	COVERAGE	NEEDED
ATTACKS/DDOS	Moderate	Moderate	Anti-DDOS	DDOS mitigation products and services are offered by a small cadre of specialized firms and MSSPs (AT&T, Arbor, Akamai). More cooperation from ISPs and infrastructure owners are needed to thwart DDOS activity.
DATA THEFT	Low	Severe	DLP, encryption	Cost and complexity of DLP still exceeds most enterprise capability. Broader adoption of whole-disk and file-based encryption, tighter integration of DLP, identity, policy silos.
EXTORTION	Low	Low	None	User education, tighter coordination with law enforcement, reputation/brand monitoring.
FENCING	Low	Moderate	Anti-fraud	Anti-fraud is mostly limited to financial services and banking. Enterprise-focused services are needed.
IDENTITY THEFT	Moderate	Severe	DLP, anti-fraud, AML	Tighter integration of DLP, SIEM, systems management and reputation monitoring.
PHISHING/ WHALING	Moderate	Severe	Anti-phishing, anti-spam, reputation/ brand monitoring	User education, reputation monitoring, improved endpoint protections.
PORN/ILLICIT	Low	Moderate	Content filtering	Broader adoption of Web content filtering.
SPAM	Moderate	Moderate	Anti-spam, antivirus	Broader adoption of anti-spam.

2.2 ANTI-MALWARE AND THE ILLUSION OF PROTECTION

If there's one common element in the enterprise cybercrime exposure grids in Figures 1 and 2, it's the centrality of sophisticated malware to enable crimes against both consumers and businesses. Indeed, modern worms, Trojan horse programs, rootkits and bots are the workhorses of the online criminal economy. They enable attackers to gain access to their target, identify new targets, survey and collect sensitive data, and carry out attacks. Alas, more than two decades after the antivirus (AV) industry sprung up in the shadows of the PC industry, the gap between threats and protection at both the consumer and enterprise levels has never been wider. To a distressing degree, modern enterprise security, as a matter of business, still hinges on the security provider's ability to know what's out there and identify it before it attacks – or after it has done its damage.

As our discussion suggests, we think that threat detection is necessary, but not sufficient to provide value to enterprises. A new model and new tools are necessary to cover the delta between what threats are out there and what current technologies protect against, as well as to help companies understand the impact of attacks after they've occurred. We'll talk about the new paradigm shortly, but first but first let's take a somewhat deeper look at how the current model broke.

2.2.1 GAMING THE SIGNATURE GAME

A variety of factors have combined to break the 'security as doctor' model. While no clear successor has emerged, we do see the protection paradigm shifting rapidly. As an example of the limitations of the security-as-doctor paradigm, consider the current market for anti-malware products. Even today, anti-malware engines are, at their root, file scanners that inspect the contents of files and compare them against a list of 'signatures' (checksums) of known threats. Anti-malware vendors have added 'signature-less' detection, bundling AV with intrusion prevention and other behavioral detection tools. They've also made great strides in disseminating signatures of new threats since the days when support techs would fax hard copies of new checksums to customers. Cloud-based 'threat intelligence services' like McAfee's Artemis, Symantec's Quorum and Trend Micro's Smart Protection Network are just a few examples.

But signature matching of known threats is still the bedrock of enterprise security – and its dirty little secret. Entrepreneurial hackers, organized criminal syndicates and state-sponsored hackers (more on the 'bad guys' later) have adapted their methods, often trivially, to take advantage of the generally favorable treatment their wares get from anti-malware products. What are some of the strategies that have been (and are) most effective at fooling current anti-malware products?

- **Mass production:** Mass production of threats has, in the past five years, proved to be the most effective weapon that online criminals have against security products that are designed to stop intrusions. Microsoft, whose visibility into threat data is unsurpassed,

recorded 116 million unique malware samples in H1, 2009.⁵ These are products of a malware industry that comprises professional and rigorous software development and scale. Experts who study malware now believe that the development resources behind its production are on par with – and in many cases, surpass – those devoted to the AV programs designed to catch them.

- **Polymorphism:** As anti-malware vendors, including Microsoft, point out, the number of unique malware programs is smaller than the amount of unique samples that AV labs receive. The population of malware samples includes loads of functionally identical software that have polymorphic features that give each copy slightly different characteristics, such as a different file size and name, or slight and meaningless changes to the internal structure of the program that, nevertheless, make it unique – the better to fool AV engines. But clearly that’s a distinction without a difference: if the malware is different enough to fool the scanner and infect your employees or users, who cares whether or not the underlying functionality or code is the same?
- **Structuring and Tuning:** Malware has long taken aim at anti-malware that might be running on target systems. Viruses frequently try to shut down or remove anti-malware executables, and a whole arms race has developed around that very challenge. However, evasion techniques targeted at anti-malware scanners and other detection tools have become even more highly evolved in recent years. In just one example, the developers behind worms like Conficker have built in features to track the IP ranges used by AV firms so they can filter out attempts to crawl Conficker-infected systems. In response, AV firms have had to enlist anonymous systems and language-specific domains to launch crawls that are able to fool the integrated filtering functions. Other firms tell us that cybercriminal groups have determined the size of anti-malware releases that elicit response from anti-malware firms, and they are purposely structuring small malware runs that slip below the radar of anti-malware research labs.
- **Quality control:** Recent years have seen malware production transform from a cottage industry to one characterized by disciplined, professional development and testing methods. Honing evasion techniques has become a cottage industry in and of itself. In the past year, researchers like Dancho Danchev have noted the appearance of hosted QA and benchmarking services⁶, including for-hire multi-engine scans that ensure released malware is undetectable by updated AV engines. These services are in addition to commodity anti-detection features that have been integrated into popular malware development toolkits that allow developers to pre-test the effectiveness of the code they release.
- **Stealth:** We also note the spread of rootkit functionality, once an obscure niche in the malware industry, to a number of popular Trojan families including Rustock, Haxdoor and others. Rootkits, which install themselves at the OS kernel level, are able to filter the results presented to applications running in user mode, effectively making resident malware invisible to both signature-based AV scanning engines and behavioral-based detection.

5. Microsoft: <http://www.microsoft.com/security/portal/Threat/SIR.aspx>

6. Dancho Danchev <http://ddanchev.blogspot.com/2009/08/managed-polymorphic-script-obfuscation.html>

- **R&D:** Profits from online crime have allowed organized cybercrime groups to fund their own R&D in order to discover zero-day exploits in common platforms such as Internet Explorer, Firefox, Adobe Reader and Flash, and to develop unique malicious programs for which no ‘threat signature’ exists.

With few exceptions⁷, there has been little recognition of the impact of the developments described above on the effectiveness of existing threat-detection tools. Industry accolades such as ICSA Labs certification and the Virus Bulletin (VB) 100 award continue to measure the accuracy of detection against historic threats, represented by the public Wild List of known malware. VB 100 certification requires 100% accuracy against this list of threats, even though that kind of test is almost without meaning for contemporary enterprises or consumers.

When tests do mix up the sample of malware, adding newly identified samples to historic threats, that 100% benchmark drops noticeably. As an example, AV Comparatives quarterly Retroactive/Proactive tests of anti-malware products attempt to measure the effectiveness of anti-malware products during a hypothetical ‘coverage gap’ between two signature updates. So engines with signatures that are current as of November 1, 2009 are run against threats that appeared between November 1 and November 7, 2009. The result? Detection rates plunge from an average of around 95% to around 55%. The top-performing product in AV Comparatives recent Retroactive/Proactive test – AVIRA – had a bare 74% detection rate. Most products were in the range of 45-55%, and a couple (Symantec, Sophos) had detection rates in the mid-30% range. Similar results accompany VB’s Reactive/Proactive test, with a majority of products detecting 50% or less of new threats and the top-performing products no more than 70% accurate against new malware.⁸

These numbers are public and have been freely available for years. Yet we’re shocked by how little change there’s been to the primacy of anti-malware as an enterprise protection tool, as we are shocked about the negligible role that hard performance and efficacy data plays in enterprise discussions about how to properly invest in security.

It’s important to remember that, so far, we’ve just been talking about mass-produced malware – commodity threats like Trojans, rootkits, viruses, worms and spyware that have been identified ‘in the wild,’ analyzed by AV research labs, labeled and tracked. We like to think of such threats as the ‘McDonald’s’ of the malware world: mass-produced, inexpensive to acquire and deploy, and ubiquitous, with low levels of differentiation within families and even across them. Anti-malware experts tell us, categorically, that this variety of threats is a big piece of what’s out there in the wild. However, they also caution that security researchers only know about a fraction of the malware that’s in circulation at any point in time, and the volume of malware production is straining the ability of research labs to process even the malware that they find⁹.

7. During her company’s 2008 analyst event, CEO Eva Chen of Trend Micro declared, openly, that the AV industry “sucks” and is falling behind malware authors.

8. RAP quadrant results <http://www.virusbtn.com/vb100/RAP/RAP-quadrant-Jun-Dec09.jpg>

9. A leading anti malware firm tells us their researchers receive 75,000 samples of “tier 1” malware – the most potentially harmful and deserving of analysis each day.

Yes, there's a long tail of low-volume or ineffective malware that fails to propagate or get noticed. Also missing are one-off designer threats that are increasingly being used by what we and others term 'advanced persistent threats,' or APTs. We'd like to spend the next section discussing what we mean by APTs and why we think they're increasingly relevant to enterprises, not just to three-letter agencies, and to the overall discussion about cyber-crime and its impact on the enterprise.

2.3 ADVANCED PERSISTENT THREATS

'Advanced persistent threats' is a term that has gained currency within government and defense circles in recent years. Increasingly, it is seen as relevant to enterprises, as well – and we think that trend will continue. APT describes a category of sophisticated, focused adversary that seeks to gain access to and maintain control over a victim's information infrastructure for the purposes of espionage and strategic or commercial advantage.

That sounds like cloak-and-dagger talk, but APTs are a huge concern among the senior IT executives we spoke with, and are particularly relevant in verticals such as finance, government, defense, critical infrastructure and high technology. Historically, APTs have been equated with state-sponsored hacking and cyberespionage, but the term applies equally well to non-state actors and profit-motivated organized online crime groups. To better understand what APTs are and how they differ from more traditional online threats, let's take a look at some recent high-profile incidents involving APTs.

2.3.1 APTS IN ACTION: AURORA, GHOSTNET

Historically, details of APTs involved in state-sponsored espionage, intellectual-property theft or large-scale fraud have been closely guarded secrets. Governments, unsurprisingly, kept mum about state-sponsored espionage directed against their networks and people, while private-sector firms worried about investor backlash or damaged reputations should they go public with details of sophisticated hacks. In the past five years, however, that wall of silence has begun to break down – possibly because the attacks have become more numerous and bold. Notably, there were the 2003 revelations about the coordinated hacks dubbed 'Titan Rain,' brought to light by Sandia National Laboratories employee and whistleblower Shawn Carpenter, and reported in the media.

In the past two years, security researchers published analysis of GhostNet, a global network of compromised hosts belonging to the Tibetan Government-in-Exile (TGIE) and other pro-Tibet organizations. That analysis was conducted by researchers at The SecDev Group and The Munk Center for International Development at the University of Toronto.¹⁰ Their analysis traced the outlines of a network of compromised hosts from Tibet-focused NGOs to those owned by the TGIE and the Office of His Holiness the Dalai Lama. Communications

¹⁰ *Tracking GhostNet: Investigating a Cyber Espionage Network*: <http://www.scribd.com/doc/13731776/Tracking-GhostNet-Investigating-a-Cyber-Espionage-Network>

from the infected hosts were monitored, and command-and-control servers for GhostNet were traced to various locations inside China. Although high-profile, GhostNet was, in many ways, a low-tech operation. Its core agents were variants of common Trojan families (Enfal, Riler and Protux),¹¹ as well as the open source (and Chinese-authored) GhOst RAT remote-access tool modified to enable data exfiltration. Similarly, researchers found they could access unsecured Web interfaces for more than one of GhostNet's control servers, making the job of surveying the network trivial. In other ways, however, GhostNet, which operated for at least three years before it was detected, comprises many of the elements that distinguish APTs from commodity malware. As researchers point out, the network itself was never very large by modern measures of botnets – just 1,295 hosts in 103 countries. But building a large, prestigious botnet was never the intention. Instead, GhostNet's operators proceeded cautiously, using high-quality, context- and target-specific phishing emails and malicious attachments (DOC and PDF) to own select systems.

Reconnaissance on email and documents uploaded from owned systems led to further targets, while providing fodder for even more pointed and effective spear phishing and social engineering of subsequent targets, and so on. Although the presence of malware was detected early on in the attack (which dates to 2002) and periodically thereafter, the attackers were able to maintain control of the host systems by keeping a low profile and through frequent updates of installed malware and the introduction of new malware and backdoors to owned systems. While the researchers at The Munk Center and SecDev allow that GhostNet may not have been politically motivated or specifically targeted at the TGIE, it's hard to see why GhostNet's operators would have kept the network itself so small, or drawn such a tight line around Tibet-related organizations if mass propagation for spam, DOS or identity theft were the true motive. And, while the attackers' use of proxies makes it impossible (or plausibly deniable) to trace the attack back to the People's Republic of China, geolocation features added to GhostNet – suggesting that the attackers were interested in the country of origin of owned systems – hint at a political element to the attacks.

Even more recently, coordinated attacks against dozens of high-profile technology and manufacturing firms – including Google, Symantec, Northrop Grumman, Juniper Networks and Adobe – reprised many of the themes seen in both GhostNet and earlier hacks such as Titan Rain. These attacks, dubbed 'Aurora,' after the name of a folder used to hold compiled malware binaries during the staging of the attack, were first revealed in mid-January¹² after researchers at Google discovered a breach of their network's security and found malware performing surveillance on select Gmail accounts belonging to political activists, as well as Google employees with access to sensitive intellectual property.

As with GhostNet, forensic analysis of the attacks pointed back to a state-sponsored actor. While an element of plausible deniability always exists, experts tell us that the provenance and forensic analysis of the Aurora malware, the locus of the command-

11. <http://isc.sans.org/diary.html?storyid=4177>

12. In a blog post by Google Chief Legal Officer David Drummond <http://googleblog.blogspot.com/2010/01/new-approach-to-china.html>

and-control servers in Taiwan and other telltale signs (e.g., the preference for using dynamic DNS to communicate between host and command-and-control infrastructure) are all consistent with the activities of China-based hackers and point to that country as the origin of the attack. The unique mixture of targets – both IP and political activists – also suggests an attacker with long-term and geopolitical interests, and one that was not primarily interested in theft of data for resale. Both the GhostNet and Aurora attacks are consistent with private-sector and government assessments of the PLA's cyberespionage and information warfare agendas, which seek to develop computer-network-exploitation techniques as a way to further intelligence-gathering and position the PLA advantageously for future conflicts.¹³

As with GhostNet, the Aurora attacks began with sophisticated spear-phishing or 'whaling' email messages sent to select Google employees. As with other attacks, the Aurora attackers had foreknowledge of their targets and were able to craft email messages that seemed to come from trusted insiders or other legitimate sources. The payloads of those messages were malicious file attachments or links to Web pages that pushed malicious code to the victim's machine. In at least a few cases, a previously unknown and unpatched remotely exploitable flaw in Microsoft's Internet Explorer browser was used in the attack to bypass defenses and push malicious code to the target system,¹⁴ though malicious files and other exploits and attacks may also have been used.

The malware downloaded in the Aurora attacks, subsequently named 'Hydraq' was a previously unknown and tailored Trojan downloader program that was undetectable by anti-malware products. While it's unclear what the provenance of Hydraq is, forensic analysis of the malware used in the Aurora attacks makes it clear that it was just one element of a preconceived and multistage operation that obviously required considerable advanced planning and modification over time.¹⁵ Once installed on the target computers, Hydraq communicated back to command-and-control servers in Taiwan and downloaded a variety of helper programs that gave the Aurora attackers complete visibility into and control over¹⁶ the victim's computer.¹⁷ Like GhostNet, the attackers moved laterally and in a methodical fashion after initial compromise, harvesting information from the victim machines to uncover other targets on Google's network with privileged access to sensitive IP and conduct surveillance on Google customers of interest. Among other things, the attackers were able to gain access to tools used by Google to respond to lawful intercept requests from law enforcement.¹⁸

13. Northrop Grumman's October 2009 report for the US-China Economic and Security Review Commission is one publicly available document that provides an excellent overview of the PLA's offensive cyber capabilities and intentions both in the military and civilian spheres.

14. As reported by SecurityFocus <http://www.securityfocus.com/brief/1063>

15. Components of the Hydraq Trojan are believed to date back as early as 2006, though it's not clear that the Aurora attack had been in the planning stages for that long, or that Hydraq was written specifically for use in Aurora.

16. Among the features Hydraq shared with the GhostNet malware was the ability to stream remote video of compromised desktops to the attackers using modified virtual network computing components. <http://www.symantec.com/connect/blogs/hydraq-vnc-connection>.

17. Press reports, citing analysis by VeriSign iDefense, initially pointed to a previously unknown vulnerability in Adobe's Reader application. However, subsequent reports pointed to a previously unknown Internet Explorer vulnerability as the first line of attack on target systems, allowing installation of a variety of Trojan downloader programs.

18. "Google attack part of widespread spying effort," http://www.pcworld.com/businesscenter/article/186786/google_attack_part_of_widespread_spying_effort.html

To be sure, there are important differences between GhostNet and Aurora. The latter attack relied on a custom and undetectable (rather than just hard-to-detect) Trojan and a hardened administrative interface, as well as encryption of inbound and outbound communications. Aurora's authors took steps (albeit basic ones) to prevent their creation from being reverse-engineered, as well. The GhostNet attackers failed to pick even that low-hanging fruit. In both cases, however, attackers were able to run highly successful operations without needing to go to extraordinary means to conceal their activities or harden their tools.

And so it goes. APTs measure success not by the technical sophistication of their attacks or their rampant proliferation (as an earlier generation of malware authors may have), but by their success at penetrating and remaining on compromised targets, as well as the amount and value of data and intelligence they can exfiltrate from those targets. As we'll see, enterprises are poorly situated and outfitted to address such threats. Giving them the tools to do so will be a major undertaking and opportunity for well-positioned firms.

2.3.2 APTs: COMMON CHARACTERISTICS

As the volume of talk about APTs increases, confusion about just what constitutes an APT has increased in proportion. In some instances, we've seen the term APT used interchangeably with malware or for specific threats such as botnets. In other cases, vendors have sought to confine APTs to state-sponsored attacks on military and government agencies.¹⁹ That's an interesting argument but one that's circular: APTs are distinguished by their focus on government networks because government entities were the first folks to start talking about APTs. It offers little in the way of developing a useful taxonomy of threats. Still others²⁰ have suggested that the term APT is misleading, suggesting 'adaptive persistent adversaries' as an alternative phrase that denotes the essentially human, rather than technological, aspect of the threat.

We think it's worthwhile, then, to harden some of the definition around APTs. One way to do that is to break down the term APT into its constituent elements.

'Advanced'

By using the word 'advanced,' we don't mean merely that the code used in APT attacks is sophisticated in nature, but that the adversaries and their methods are advanced. Indeed, much of the malware used in these attacks overlaps or shares code with existing families of 'commodity' malware that is in broad circulation. Given the sophistication evident at the high end of that market (we note the recent analysis of evasion techniques built into the newest generation of the TDL3 rootkit family as one example²¹), as well as the rela-

19. Will Gragido at Cassandra Security, notably. <http://cassandrasecurity.com/?cat=78>

20. Notably Nick Selby and Scott Crawford in a ThreatPost column: http://threatpost.com/en_us/blogs/its-adversaries-who-are-advanced-and-persistent-012610

21. Rootkit.com <http://rootkit.com/newsread.php?newsid=979>

tive affordability of custom malware from coders for hire²², it would be difficult to argue, in blanket fashion, that malicious programs used by APTs are more sophisticated than commodity malware.

Looking at the details of publicly disclosed APT attacks like GhostNet, Aurora, Titan Rain and the recently publicized attacks on Western energy firms, the term ‘advanced’ better applies to the planning and execution of the attack, as well as the resources available to the attacker, than to the malicious software used. In particular, the adversary’s reconnoitering of its target, the fashioning of sophisticated, credible social-engineering attacks that allow them to bypass local protections, and the methodical, deliberate and stealthy nature of their activities subsequent to gaining access all contribute to the designation of an ‘advanced’ threat.

Less discerning adversaries may target widespread and easy-to-exploit vulnerabilities, such as those found in versions of Microsoft Windows or in common applications such as Adobe’s Reader or popular Web browsers, or they may use email, IM and the Web as attack vectors. So, too, APTs. But when remote attacks via email, IM or the Web are used, they are likely coupled with a targeted (aka ‘spear’) phishing or social-engineering attack against high-value targets (aka ‘whaling’) after considerable online reconnaissance. (Social networks and the Web now make such open source intelligence gathering easy and powerful.)

APTs may also leverage a remote exploit of an OS, browser or application vulnerability to gain a foothold on a target host or network. However, they might just as easily be introduced by a trusted insider or outsider (contractor, supply-chain partner) with legitimate credentials, which are then leveraged in later stages of the attack. With a boundless sea of potential commercial targets, most profit-motivated hackers, even sophisticated ones, are unlikely to expend the same amount of energy on a single target.

‘Persistent’

As for the other distinguishing characteristic of APTs – their ‘persistence’ – once again, the idea behind this is to note a qualitative difference with less discriminating attacks and adversaries. By persistence, we note the focused and dogged nature of APTs. APT adversaries know, often in great detail, what information (or at least what types of information) they are after before their attack begins, and are single-minded in their pursuit of it throughout that attack. The failure of any single foray is unlikely to end a mission. Rather, adaptive persistent adversaries will have researched a number of attacks to be tried in parallel, or serially, to achieve their objective. If targeting high-level officers fails, advanced adversaries will go after smaller fry, then move toward their desired target(s). If remote compromise is unlikely, APTs may consider flipping an insider or planting one of their own within the target organization.

22. Dancho Danchev features an intriguing look at the menu of choices (with prices included) of a malware for hire author on his blog. <http://ddanchev.blogspot.com/2008/07/coding-spyware-and-malware-for-hire.html>

'Threat'

The use of the term 'threat' has generated much confusion – leading some within the security community to equate APTs with malware. As we noted above, while APTs may well use software exploits, viruses, worms, bots, rootkits, Trojan horse programs or other malware as part of their attacks, APTs are adversaries, not simply software tools. However, it is worth observing the obvious – that APTs frequently rely on malware, and that there are distinctions worth noting between the kinds of malware used in APT attacks and in non-APT attacks.

At the ten-thousand-foot level, we believe that the distinctions between APT-associated malware and what we call 'commodity' malware are usually a matter of degrees rather than black-and-white differences. Feature-wise, the tools used by APTs are similar or identical to non-APT-associated malware. They may, in fact, borrow code from commodity threats. But APTs privilege persistence over propagation. Remote access to target hosts, networks and the sensitive data they contain is the ultimate objective of all kinds of malware. However, where commodity threats are likely to be loud and sloppy, APTs will be silent and deliberate – harvesting data on compromised hosts carefully before moving on, laterally, to additional targets. Initial compromise may come by way of vulnerabilities in common applications in both cases, but commodity threats will be more likely to use publicly disclosed holes and exploit the patch window. APTs will be more likely to rely on privately researched and previously unknown vulnerabilities or custom exploits, specific to the target network and system.²³

Communications with external command-and-control servers may be a common feature of both APT and non-APT attacks, but APTs may take extra steps to disguise that communication – using command-and-control servers on the same subnet, or from legitimate-seeming domains and IP addresses to avoid detection blacklists and other monitoring tools. Communications to and from command-and-control networks will piggyback on top of legitimate traffic (HTTP, etc.) and pass through heavily used ports as much as possible in order to make the 'performance/security' trade-off of filtering them unpalatable.

As an example, the Conficker worm is a sophisticated piece of malicious software. It cloaks itself on Windows systems, uses code-hardening to resist reverse-engineering, pushes down daily binary updates to complicate detection, removes common anti-malware products on infected hosts, and patches the MS08-067 Windows server service vulnerability that the worm exploits to infect hosts in order to avoid later detection by patch management tools.²⁴ However, the worm's primary purpose is propagation and the creation of follow-on services (spam, DOS) that leverage its network. As a result, the worm exhibits telltale signs that make it relatively easy to distinguish. (SRI International-

23. *An analysis of the Hydraq Trojan used in the Google Aurora attacks noted just this – the threat, itself, was unique, but had components that were not. Most notable, in terms of sophistication, was the zero-day exploit used to gain access to hosts.* <http://www.symantec.com/connect/blogs/trojanhydraq-typhoon-teacup>

24. *"An analysis of Conficker's Logic and Rendezvous Points," SRI International.* <http://mtc.sri.com/Conficker/>

al's analysis notes chatting on ports 53, 80 and 445, as well as periodic spikes in DNS activity as the Conficker drones poll the Conficker command-and-control network for binary updates.)

Like most commodity malware, Conficker is playing a numbers game²⁵, and is counting on the long tail of organizations that lack such expertise for its targets. In contrast, threats like GhostNet and Hydraq, or even the 'blabla' malware used by the Gonzalez crew, took extra steps to fly under the radar and avoid detection – previously unknown vulnerabilities and exploit code, custom malware, or some combination of those. In the end, those simple steps were enough to earn the threats the label 'sophisticated' and to maintain their hold on the target network for weeks, months or longer.

FIGURE 3: FEATURE CHECKLIST – COMMODITY MALWARE VS. APT

	APT	COMMODITY MALWARE
ROOTKIT	Yes	Maybe
POINT OF ATTACK	Spear phishing, Web, rogue/compromised insider, removable device, wireless	Spam, phishing, Web
PURPOSE	Industrial and state-sponsored espionage, cyberwarfare, large-scale financial fraud	Spam, DOS, small-scale financial fraud
EXPLOITS	Low and slow; leverages private exploits whenever possible against applications and infrastructure known to be deployed on target network.	Indiscriminate and using known exploits of common platforms and applications; goal is to build a large network for premium services (DOS, spam, data theft)
PROPAGATION	Careful, targeted at known high-value assets and individuals	Indiscriminate, leveraging platforms (email, social networking, etc.) that provide maximum reach with minimum effort
PLATFORM	Windows, Unix, Linux, mainframe, SCADA, Mac, mobile (BlackBerry, Windows Mobile, iPhone, etc.)	Windows
COMMAND AND CONTROL	Restrained, focused on data exfiltration; uses common channels (HTTP) or uncommon ones – IPV6, etc.	Chatty; fast flux; focused on propagation
TARGET	Government (military and civilian branches), finance, IP-sensitive verticals (defense, tech, biotech, pharma, industrial)	Consumer, SMB, retail and merchant banking, e-commerce

25. SRI's census noted more than 10 million Conficker IP addresses at one point. <http://mtc.sri.com/Conficker/#appendix-1>

2.4 CYBERWARFARE AND ATTACKS ON CRITICAL INFRASTRUCTURE

Another trend that is upsetting the security status quo is the rising specter of cyberterrorism and cyberwarfare carried out between states or by non-state actors. While cyberwarfare and attacks against critical infrastructure have long been postulated (and perhaps conducted, covertly), events in recent years have underscored the degree to which physical and cyber infrastructures have become intertwined.²⁶ At the same time, events have shown that offensive cyberattacks are now a weapon in the arsenals of advanced and developing nation-states, including Russia (as evidenced by that country's sustained cyber and physical offensive against former satellites Georgia and Estonia²⁷), China, Israel, India and others.

News reports in recent years have surfaced intelligence reports about successful attacks on power grids outside the US²⁸, as well as attacks on US military and research facilities and the theft of military secrets (such as the US's advanced Joint Strike Fighter Jet²⁹) and other classified information within the US. A partial list of critical-infrastructure attacks and incidents maintained by Industrial Defender³⁰ and dating back to 1982 notes around 60 discrete recorded events, the bulk of them occurring since 2007. It's likely that a complete list of such incidents – both classified and unclassified – would stretch back further and contain far more data points.

For the IT security sector, greater emphasis on cyberwarfare and attacks on critical infrastructure presents both challenges and opportunities. First, billions of dollars in federal investment to help secure the nation's critical infrastructure will ensure that the IT security pump stays primed, even during the lean years ahead as the US economy climbs out of the deep hole it fell into beginning in 2007. Various federal initiatives are already creating demand for cyber-security and digital-forensics training, malware and threat analysis, and threat intelligence services. Recent years have already seen numerous M&A deals in this area, including *Raytheon's purchase of Oakley*, as well as McAfee's purchase of supervisory control and data acquisition (SCADA) security specialist *Solidcore* and UK firm QinetiQ's *purchase of cyberintelligence and monitoring firm Cyveillance*.

26. *The outbreak of the SQL Slammer worm in 2003 did disrupt systems within the electric distribution system, though no known power outages were linked directly to the worm. (Claims that the massive Northeast blackout was a byproduct of the Slammer outbreak have been refuted.)* www.esisac.com/publicdocs/SQL_Slammer_2003.pdf

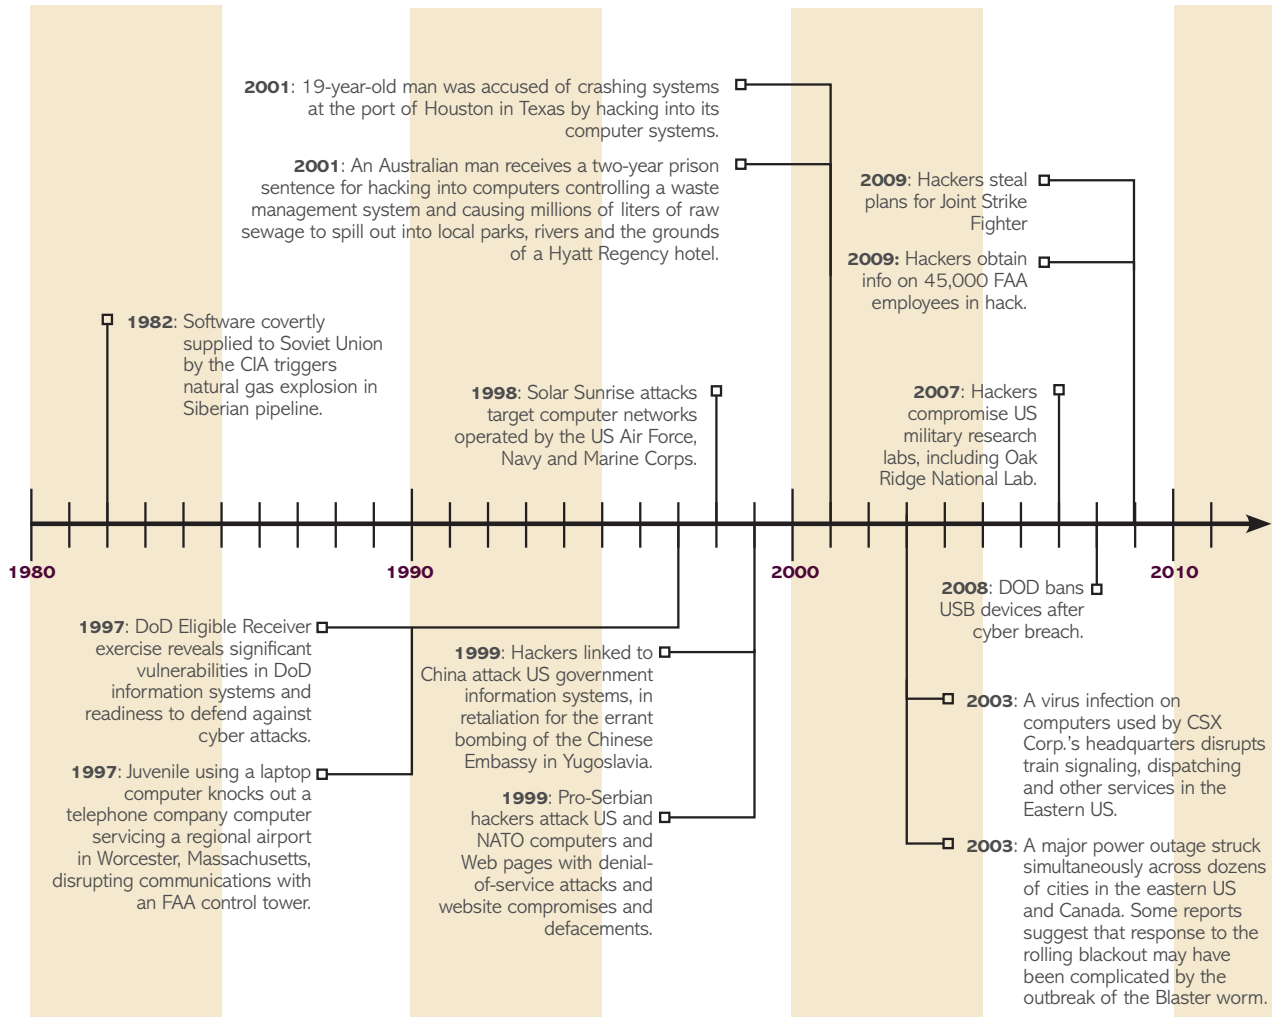
27. <http://bits.blogs.nytimes.com/2008/08/11/georgia-takes-a-beating-in-the-cyberwar-with-russia/> and <http://news.bbc.co.uk/2/hi/europe/6665145.stm>

28. <http://www.sfgate.com/cgi-bin/article.cgi?f=/n/a/2008/01/18/national/w122440S64.DTL>

29. <http://online.wsj.com/article/SB124027491029837401.html>

30. <http://www.industrialdefender.com/news/incidents.php>

FIGURE 4: CRITICAL INFRASTRUCTURE ATTACKS



Source: Industrial Defender

But SCADA and other platforms used to manage and operate critical infrastructure challenge the Windows-centric world of enterprise security, forcing greater breadth of support for the specialized platforms (including older Windows variants) that SCADA software often relies on. Beyond that, critical-infrastructure providers operate with an entirely different set of priorities than enterprises do. As Brian Ahern, President and CEO of Industrial Defender, pointed out in a conversation with The 451 Group, enterprises prioritize the confidentiality and integrity of their data over availability of the data. Critical-infrastructure providers, however, put availability above every other priority, which means that mission-critical assets have to operate 24/7 without interruption. Such a requirement is out of scope for many existing enterprise security tools, which assume that small-percentage drops in utilization or short lapses in availability will be tolerated in the name of security.

In short, securing public and private entities from attacks sponsored by nation-states or non-state actors presents a huge opportunity for security firms of many different stripes. But playing in the critical-infrastructure market will require firms that have focused solely on protecting traditional LAN and WAN environments to broaden both the tools and services they offer in order to be relevant to critical-infrastructure providers. We take a look at what changes that might entail later in the report.

SECTION 3

FUD Factor: Hype, Compliance and the Problem of Information Asymmetry

As we discussed in Section 2, cybercrime and APTs present a number of challenges to enterprises and call into question the utility of legacy IT security investments. Professional adversaries motivated by profit and politics, low- and slow-moving attacks, multivector threats that leverage the Web, zero-day vulnerabilities and ‘layer 8’ – IT users themselves – challenge single-function security devices and threat-based protections that are reactive rather than proactive.

At the same time, the IT security industry’s response to the evolution of threats – a proliferation of specialized security appliances and agents – has created a cost-and-complexity crisis within organizations that must be resolved. Enterprise DMZs are already a forest of expensive single-use appliances, while desktops slow down under the demands of loosely integrated endpoint security suites, patch and configuration management agents, and other specialized tools. As new threats come along – spam, spyware, rootkits, botnets, Web-based attacks, etc. – new products are invented to address them.

But are the solutions being offered to address these new threat vectors likely to be effective, or are they just expensive Band-Aids that will force attacks down a different, yet unseen avenue without offering any net improvement in security? These are difficult questions to answer. There’s consensus that the Internet is a – if not *the* – critical infrastructure of our 21st century economy, and that cybercrime, cyberterrorism and cyberwarfare are new fronts in a war that must be fought. But, as Aeschylus famously observed: “In war, truth is the first casualty.” While we might agree on the broad outlines of the problem, there’s hardly any consensus within either the private or public sectors on the specifics of the threat or the proper means with which to address it.

One of the most bedeviling issues confronting those trying to raise awareness about the problem of cybercrime or related phenomena like cyberterrorism and cyberwarfare is sizing the problem in a way that is meaningful to those with the power to address it (CISOs, industry and political leaders, or regulators) and that is also empirical and proportional to the problem. Alas, the public discussion about these problems lacks most of these qualities. It tends to be sensational, driven by news stories and headlines about data breaches at marquee institutions, and pointillistic – focusing on discrete hacks or compliance objectives, but missing the opportunity for a broader discussion about preparedness, prevention and provable security. Stories about hacks at firms such as Citibank, RBS and TJX and others paint a picture of large, sophisticated financial firms at the mercy of shadowy crime syndicates, but rarely isolate specific causes or remedies that could have prevented attacks (probably because that’s nearly impossible to do).

Finally, despite years of dire warnings, advocates of better computer security have few ‘cyber Enrons’ to point to – instances of gross mismanagement or lax oversight that

resulted in spectacular failures of seemingly healthy enterprises. With the exception of very specialized firms like credit card processor Card Systems Inc, companies that have been the targets of some of the largest cybercrimes – ChoicePoint, TJX, Hannaford Supermarkets, OfficeMax and even Heartland Payment Systems – continue to operate, and some have even prospered in the wake of their breach announcements. For most firms, in other words, cybercrime is better characterized as a manageable (albeit growing) business risk, rather than an existential threat.

In this section, we'd like to talk briefly about what we see as a major obstacle to addressing the problems of cybercrime and state-sponsored attacks: namely, the dearth of accurate information about the nature and dimensions of the threats facing organizations. We would also like to make an argument for changes that will result in a more reality-based discussion of online threats, including organized cybercrime, cyberterror and state-sponsored espionage and cyberwarfare. The goal is for organizations to direct their sparse resources toward programs and technologies that will yield dividends – an improved security posture, fewer successful attacks and lower overall costs for security.

3.1 SIZING THE CYBERCRIME AND APT PROBLEM

How big is the cybercrime problem? One of the challenges is defining the term 'cyber-crime.' Taking a narrow definition of just Internet-related crime, there are estimates such as the one from the FBI's Electronic Crime Complaint Center. That organization put the total dollar losses attributable to Internet crime for 2008 at just \$265m, or around \$931 per complaint. There are other consumer-focused surveys, such as Consumer Reports' State of the Net, which put the dollar amount attributable to various types of online crime (phishing, viruses, spam and spyware) at around \$8bn in 2009.³¹ A widely cited report from McAfee and Purdue University's Center for Education and Research in Information Assurance and Security put the cost of security breaches at \$4.6m per firm at the 1,000 global firms the group surveyed, measured by the value of lost intellectual property. That figure puts total estimates for losses attributable directly and indirectly to online crime at around \$1 trillion globally.³²

That's a wide range of estimates that varies depending on who is being measured (consumers and SMBs vs. large corporations) and what is being measured (online scams and schemes vs. targeted theft of intellectual property). When researchers do crawl behind the numbers, a different picture often emerges. In just one example, Microsoft researchers Cormac Herley and Dinei Florencio used data from an Internet Explorer anti-phishing toolbar to look at incidents of phishing and estimate dollar losses attributable to them. Based on about 500,000 user encounters with phishing sites, the two discovered that just a small fraction of users, around .37%, end up handing over credentials to the fraudsters. Based on that figure, Herley and Florencio estimated that losses attribut-

31. President Obama cited the Consumer Reports data in a May 2009 speech on the importance of securing the nation's cyber infrastructure.

32. McAfee Inc. <http://resources.mcafee.com/content/NAUnsecuredEconomiesReport>


able to phishing were in the neighborhood of \$60m annually, compared to surveys by analyst firm Gartner that put phishing-related losses at \$3.2bn (based on a 2007 survey of around 4,500 US adults). Rather than a lucrative online racket, Herley and fellow researchers theorized that phishing is a low-skill, low-profit operation characterized by ‘over-grazing’ – too many phishers chasing after too few phish, so to speak.³³

A similar problem emerges even when looking at a single incident of cybercrime, let alone an entire sector, national or global economy. As an example, consider the breach at Massachusetts retailer TJX. Measured by stolen credentials, the TJX breach was massive, involving the theft of account information on around 45 million credit card accounts. But what were the costs to the company to recover? It turns out to be a much harder question to answer. The data security firm IPLocks estimated the breach would cost TJX \$4.5bn (at a rate of \$100 per stolen record). The Ponemon Institute put the number even higher: at \$8.6bn (using a figure of just over \$180 per stolen record). Analyst firm Forrester put the cost at between \$90 and \$305 per stolen record, with total breach costs varying accordingly.

Which of those numbers is accurate? One way to determine an answer is to look at how much money TJX has allocated to recovery costs. To date, TJX has set aside \$170m to pay for breach-related expenses. Settlements with most parties affected by the hack – including banks and state attorneys general – is closer to \$75m,³⁴ or around \$1.60 per record. (Much of the money is in the form of ‘set asides’ to credit card issuers for the cost of card replacements and fraud claims that may or may not be filed.) There are other costs that should be included, but which are harder to come by: the cost of network assessments during and after the breach, additional investments in IT staff and infrastructure to respond to the findings of those assessments, and fines. In addition, TJX suffered reputational damage that will be hard to erase, though the company’s health, as measured by its stock price and financial performance, appear to have been little affected by the breach since it was first reported in 2007. (Given the complex ramifications of identity theft, we’re betting that consumers whose identities were stolen as a result of the breach may find it harder to have their financial reputations restored than TJX did.)

33. Some of this data was included in Herley’s Black Hat presentation on “Economics and the Underground Economy.” <http://research.microsoft.com/en-us/people/cormac/>

34. “TJX agrees to settle another breach lawsuit for \$525,000,” *Computerworld*, http://www.computerworld.com/s/article/9137491/TJX_agrees_to_settle_another_breach_lawsuit_for_525_000

FIGURE 5: ESTIMATED LOSSES FROM TJX DATA BREACH (\$B)

Why is it important to try to accurately size the cybercrime and APT problem? While speculation about the size of the cybercrime problem may be understandable, we believe that it also leads to misplaced investments or, even worse, inertia. As Microsoft's Herley has noted, overwhelming statistics about the dimension of the online crime problem tends to divert resources from where they're needed. It can also instill a sense of hopelessness in the minds of consumers and business owners about their ability to take (inexpensive) corrective actions that may limit their exposure. As an example, Herley suggested that a public information campaign based on hard, impartial data on, say, the advantages of strong authentication might do far more to promote safe behavior than scary statistics about the ubiquity of identity theft.³⁵

Alas, the desire for reliable metrics on the cybercrime problem and the smaller but no less serious threat of state-sponsored hacking is likely to remain unfulfilled. In North America, large swaths of the critical IT infrastructure lies in private hands, where owners have no requirement or incentive (and lots of disincentives) to disclose breach information. Were such requirements to exist, compliance with them would be impossible to measure, and any statistics would be presumed to be unreliable. In the case of state-sponsored hacking, responsibility for such attacks is often hard to prove, while governments are understandably loath to acknowledge what they know of such activities for a range of reasons. Finally, among those entities that do aggregate attack data – most security software and infrastructure companies, telecommunications providers and the like – that data is increasingly bent to serve the marketing needs within those firms. Even seemingly comprehensive and unbiased reports from private-sector firms³⁶ have to be

35. A more recent paper by Herley argues that ignoring advice on safe computing practices is economically justifiable, given that the cost of added protections outweighs the cost of attacks.

36. Microsoft's semiannual Security Intelligence Report and Verizon's annual Data Breach Study are two good examples.

viewed through the lens of that company's marketing needs and objectives. Too often, however, we find that the conclusions of such reports are accepted *prima facie* and, in the absence of countervailing data from unbiased sources, become the final word on both cybercrime threats and possible responses to those threats.

3.2 WHACK-A-MOLE PAYS THE BILLS

Absent reliable data on cybercrime trends from impartial sources, our understanding of problems like cybercrime and state-sponsored hacking are highly susceptible to spin by interested parties to support a marketing (or political) objective. This phenomenon isn't new. The AV industry learned long ago that news about threats, rather than undermining support for their products, merely stokes demand for more and better threat protection. Or, as an executive at one (highly respectable) IT security firm told us, "Whack-a-mole pays the bills."

In part, the 'whack a mole' dynamic is unavoidable in the fight against computer crime. As in any other area of criminal behavior, cybercriminals prosper by staying one step ahead of both their victims and law enforcement. The practical consequence of this is a steady stream of new attacks, new methods of exploitation and new forms of fraud. 'Good guys,' broadly defined, constantly have to adapt their protections to the new threats. And as soon as protection comes in line with threats, new threats are developed.

The last decade has been a case study in that dynamic, with threats and attack vectors evolving at a steady clip to stay in front of defenses – from spam to spyware to botnets and sophisticated rootkits and Trojans. At each step of the evolution, criminals found ways to identify and embrace new or just ubiquitous technologies and leverage them for illicit gain. And, at each stage along the way, protection lagged threats as security incumbents wrestled with the implications of new threats and attack modes. Funny as it seems now, there was a heated debate within the anti-malware community about whether spam was a problem that warranted equal attention with viruses and other kinds of malicious code.³⁷ Similarly, the last decade has seen established security vendors move slowly to address the threat posed by spyware, and by a new generation of kernel-mode rootkits capable of evading detection by signature-based virus scanners, despite researchers' warnings about their increasing prevalence. As one top security researcher at an anti-malware firm responded, candidly, when asked about the silence on the rootkit problem, "Why would I talk about a problem that we can't solve?"

And, at least in the short term, that approach appears to be working. Security vendors have succeeded in selling enterprises an increasingly complex menu of security products that now includes technologies as diverse as gateway and desktop AV, anti-spam and firewalls, application firewalls, intrusion detection, application control, data-leak prevention (DLP), data encryption and security incident management. The ever-shifting

37. After much debate, the Virus Bulletin Conference allowed presentations on spam beginning in 2003. <http://www.virusbtn.com/conference/vb2003/programme/index>

threat landscape has also proved to be a fertile ground for the development of targeted security wares, providing an incentive for the continued attentions of the venture capital community. As we and others have noted, the percentage of IT budgets going to security has increased dramatically in recent years.³⁸

3.3 FUD AND ITS DISCONTENTS

While the evolution of threats is clearly one (legitimate) driver of IT security investment, there's no debating that vendor-backed fear, uncertainty and doubt (FUD) is another. From viruses to spyware to banking Trojans, botnets and APTs, IT security companies have found it easy enough to stoke both customer and popular fears of online threats in order to advance their marketing and sales objectives. A compliant press and the general lack of technical understanding of online threats and crime make it easy enough to spin gold from the straw of online malicious activity.

Examples of this are almost too numerous to count, and are as old as the computer security industry itself. To take a recent example of the FUD game, consider the media maelstrom that erupted in early 2009 over the Conficker worm and its supposed 'April Fool's' surprise. If you recall, the story then was about the millions of worm-infected machines that were lying dormant, but would suddenly respond to secret commands, wake up and do something... most likely something bad. After prime-time news magazine *60 Minutes*³⁹ picked up the story and ran with it, a full-on scrum ensued with wall-to-wall Y2K-style speculation about the dark purpose of Conficker's authors and the disruptions it might cause when it 'goes off.'

The *60 Minutes* story was good PR for the AV industry, especially Symantec, which got to show off its experts, research and SOC. Unfortunately, in the hands of the mainstream media, the story quickly became about the April 1 Conficker 'doomsday,' rather than the less sensational but more weighty discussion of the botnet and malware problem, the links between malware and problems like spam and identity theft, and so on. When that doomsday failed to materialize, it was natural to conclude that the security industry was, once again, crying wolf.

Even more recently, a front-page story in *The Wall Street Journal*⁴⁰ about data theft at US corporations subsequent to infection by the Zeus Trojan was driven by a report from security firm NetWitness. As with Conficker, the threat posed by Zeus, an information-stealing Trojan, is real enough. Both warrant attention from the press, lawmakers, regulators and the IT community. The problem is that when the discussion is driven by news that's ginned up by the marketing divisions of IT security firms, the picture of the overall threat that emerges is pointillistic rather than blended. Zeus and the Kneber criminal group that use it are, by no means, unique or particularly new. Nor is the

38. <http://www.darkreading.com/security/management/showArticle.jhtml?articleID=212700661>

39. <http://www.cbsnews.com/stories/2009/03/27/60minutes/main4897053.shtml>

40. <http://online.wsj.com/article/SB10001424052748704398804575071103834150536.html>

theft of enterprise data by Trojan-infected hosts. In fact, both trends were well documented before the NetWitness study came out.⁴¹⁴¹ Rather than provide that context, however, or delve deeper into larger issues, the Zeus story, in its various incarnations, spun a single instance of an endemic problem as an example of a ‘new’ threat akin to the Google Aurora hacks. (‘Broad New Hacking Attack Detected,’ the headline read for the WSJ Zeus story.) True, both attacks involved the theft of corporate data, but the targeted Aurora compromise that was focused on intellectual property theft and political targets is utterly different from the wholesale Zeus compromises, which are financially motivated and less discriminate.

We can think of a lot of important and useful lessons that might have been derived from a balanced discussion of the endemic Zeus malware: the sophistication of the newest generation of malware, the fact that the products sold by the anti-malware incumbents commenting on the infection aren’t capable of keeping up with variations in the Zeus family and preventing compromises, the dangers posed by corporate data theft, and so on. As it ended up, this vendor-driven story merely stoked incoherent fears, obscured some important context on the history of the threat and failed to propose remedies.

Why does FUD matter? We believe that, while unavoidable, the centrality of vendor-driven FUD to the ongoing discussions of cybersecurity makes collective action to address threats harder, not easier. Sensational stories may drive business and discussion in ways that benefit individual players, or even entire sectors. But they also factionalize the IT security community in ways that are counterproductive. This dynamic emerged, clearly, with the Zeus Trojan story – with larger firms like Symantec and McAfee pooh-poohing the findings of the NetWitness research, creating (in this case, justifiable) confusion among customers and within the larger business community about the seriousness of the threat itself. In fact, as we discuss in our next section, the ongoing use of FUD has led to a crisis of confidence between security vendors and their customers that already threatens the ability of the IT community to respond to a fast-changing threat landscape.

3.4 THE PROBLEM OF INFORMATION ASYMMETRY

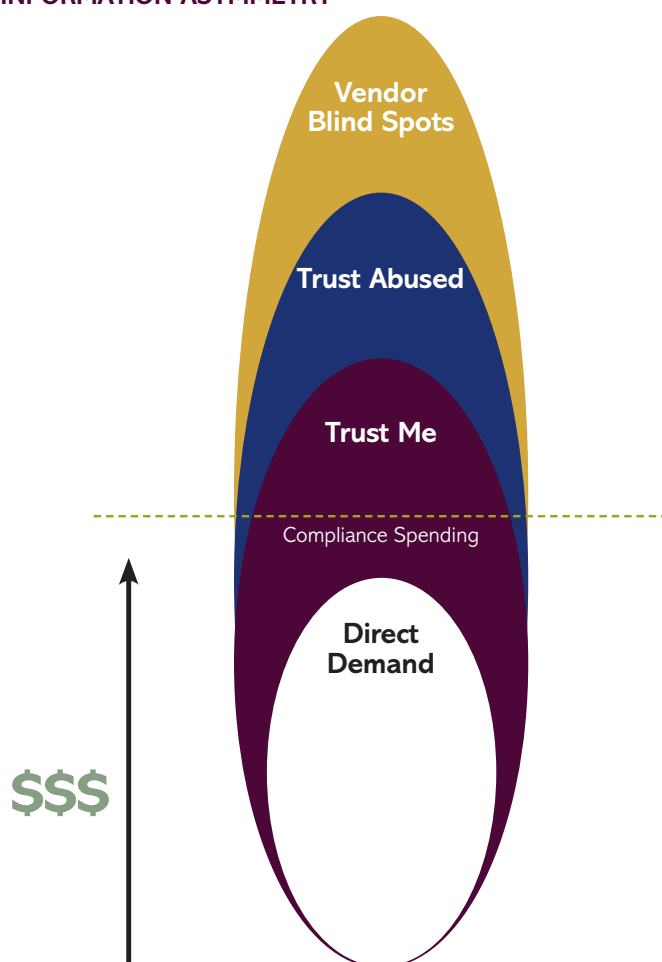
While it’s possible that all this FUD translates into better enterprise security, anecdotal evidence and surveys of IT leaders suggest otherwise. In fact, confidence in the efficacy of existing IT security investments is quite low among the enterprise IT leaders we spoke with for this report. They lack confidence in the ability of their current tools to keep pace with the latest malicious programs and attacks (Web- and application-based attacks are of particular concern). They are also doubtful of their ability, given their current tool set, to spot rogue insiders, information-stealing malware or bots, or other sophisticated threats. Most of all, they worry about ‘black

41. <http://www.krebsonsecurity.com/2010/02/zeus-a-virus-known-as-botnet/>

swans’ – the next paradigm-shifting threat or attack (think SQL Slammer) that will catch everyone unawares and against which they will have little forewarning or protection. Given recent history, we think these concerns are well founded. We also think the IT security industry bears the weight of some responsibility for the generally low opinions and flagging optimism of its customers.

As 451 Enterprise Security Research Director Joshua Corman argued in his recent report on ‘information asymmetry,’ security vendors have been on a slippery slope with the truth for most of the last decade. As the number and complexity of threats outpaced the ability of simple tools to counter them (virus and antivirus, as an example), security vendors have been put in the role of security advisor, rather than merely security provider: educating customers about new threats before addressing the solution to that threat. As we’ve noted, however, that role has proved to be subject to abuse, as simple solutions fall short of the mark and vendors fall into the role of bending their advice to suit the capabilities of their product offerings and, eventually, to talking only about the subset of known threats for which they have solutions – steering clear of those known threats that they can’t easily address, and ignorant of still other developing attacks and threats. Protection, then, is for a subset of a subset of the threats that are actually in the wild.

FIGURE 6: INFORMATION ASYMMETRY



As an example, consider the ways in which AV vendors have morphed from sellers of a ‘fix’ for boot-sector and macro viruses into trusted enterprise resources on enterprise threats of all types. In the meantime, these vendors have consistently been caught flat-footed when paradigm shifts occurred and new threats arose. They (arguably) soft-pedaled the relevance of threats like spam, spyware, rootkits and botnets to their customers until those ‘markets’ had proved themselves or they had ready offerings in those product categories. That dynamic is true even today, as these incumbents find themselves in the untenable position of having to advise customers about the dangers of next-generation threats while also arguing for continued investment in legacy products that are incapable of addressing the threats of which they’re warning.

3.5 A COMPLIANCE CONUNDRUM

While vendor-driven FUD has long been with us, regulatory compliance has, within the last decade, become an even more potent force that’s driving the increase in IT security expenditures. Evolving compliance mandates in the US and elsewhere such as PCI DSS, HIPAA and the EU’s Data Privacy Directive, not to mention an array of federal initiatives and state data-privacy laws in the US, have stoked demand for technologies as diverse as AV software, vulnerability-scanning tools, full-disk encryption and Web application firewalls.

We think that, insofar as compliance mandates reinforce security best practices, they are enormously effective. Our research has, at times, been critical of prescriptive regulations such as PCI DSS, but there’s no doubting that toothy regulations like PCI have utterly transformed the enterprise security landscape and, especially among midtier and smaller merchants, have directed resources toward IT and data security that might have otherwise gone elsewhere. However, it’s also clear that the growing centrality of regulatory compliance to enterprises exacerbates the problems of FUD and information asymmetry. Faced with the inevitability of an audit but only the possibility of a malicious attack, enterprises direct spending and resources toward controls that will satisfy their compliance needs, regardless of whether they improve overall security posture.

The more prescriptive the regulations and the more force behind them, the more they serve to retard innovation and experimentation that might, ultimately, lead to better security. As an example, the PCI DSS mandates specific technologies such as AV software, desktop firewalls and Web application firewalls.⁴² That is because, in large part, those technologies represented the state of the art and best practices at the time the PCI DSS regulations were being written. The problem comes when those mandated technologies lose their effectiveness, as it can be argued that AV software already has. With two-year cycles for updates to PCI DSS and the accretive nature of most regulations, it is unlikely that the AV mandate will disappear anytime soon, even as attackers move nimbly around existing protections, and alternatives to multifunction anti-malware

42. PCI 6.6 allows Web application vulnerability testing and remediation in lieu of a WAF.

suites and signature-based detection proliferate. (See Section 4 for more on this.) With AV a ‘must have,’ enterprises are, more or less, discouraged from swapping out ineffective protections for new ones – any new detection tool will need to run alongside the mandated AV, even if it is technically a replacement for it.

We’ve already seen instances of high-profile firms that followed compliance mandates to the letter, but still found themselves victimized by attackers. Heartland Payment Systems is the most prominent example here. The recent Aurora attacks, as well as news about the continued success of criminal groups behind the Zeus Trojan, also underscore the degree to which PCI-mandated protections are falling short – codifying legacy tools and controls (malware scanning, vulnerability scanning, firewalls) while failing to adapt to new threats and attacks (designer malware, zero-day vulnerability exploits, Web-based attacks). In fact, given the amount of investment in compliance, there is a strong disincentive to modify compliance guidelines once they’ve been established – especially if that would mean swapping out one (more effective) technology for a technology that organizations have been mandated to acquire. We see that dynamic playing out now with PCI DSS, with organizers already promising ‘no major changes’ in requirements for the next version of the standard,⁴³ despite a threat landscape that is changing rapidly.

And as threats become more sophisticated and difficult to track, attention is increasingly directed toward achieving compliance, regardless of outcome. Unable to properly define the threats or ‘security’ in a way that’s meaningful, in other words, compliance has become a stand-in for security, even though the gaps between one and the other are all too plain to see.⁴⁴

For ISVs that have products positioned in the right category, regulations can be the goose that laid the golden egg. For organizations that are looking to rationalize and align their security investments and infrastructure, however, regulatory mandates are often received as unwelcome and poorly scoped injunctions from faceless state, federal or industry regulators.⁴⁵ And by no means is this process at its end. In just one example, there are multiple pieces of legislation under consideration by the US Congress that could expand IT security regulations in still new directions, mandating heretofore elective functions such as network- and application-penetration testing by government agencies or businesses.⁴⁶ As is often the case, companies with an interest in seeing that kind of analysis mandated are tapping connected insiders and working the halls of Congress to get their technology mandated by law. While the addition of such a requirement may be warranted, the cumulative effect of these requirements – new technologies

43. “No Major PCI Revision Expected in 2010” http://searchsecurity.techtarget.com/news/article/0,289142,sid14_gci1379760,00.html

44. Heartland Payment Systems had passed a PCI audit during the period when a packet sniffer operated by Albert Gonzalez was collecting credit card data from its payment processing network.

45. Consider the difficulty Massachusetts regulators have had implementing a poorly scoped data privacy law. Implementation of that law (set to take effect on Jan 1, 2010) was delayed twice and trimmed back as businesses in the state complained about the difficulty and cost of implementing data-encryption requirements and verifying the security of third parties. http://www.mass.gov/?pageID=ocapressrelease&L=1&L0=Home&sid=Eoca&b=pressrelease&f=20090212_idtheft&csid=Eoca

46. The US Cybersecurity Act of 2009 and US Info and Communications Enhancement (ICE) Act of 2009 are just two.

are added to the list, but old ones are never removed – will be to make IT security sclerotic and less effective in the face of criminal enterprises and state-sponsored espionage of the kind we’ve discussed.

The challenge for regulators in both the private and public sectors, then, is to find a way to encourage – indeed, require – best practices around IT security, but to also maintain agility and flexibility in the system to adapt to a changing threat environment. At their best, regulations force the hands of reluctant enterprises that, left to their own devices, would give short shrift to IT security. They can also put the spotlight on critical but under-addressed areas (such as data security). They can help establish the parameters by which organizations can evaluate the effectiveness of their existing IT security investments and identify areas that require more attention. At their worst, prescriptive regulations like PCI DSS cement legacy technologies (AV, firewall, IPS, etc.) in place and constrain the ability of IT decision-makers to make fundamental changes to their IT security infrastructure in light of an objective assessment of the threats and exposures facing their business.

Despite the current single-minded focus on compliance, most every regulation represents a low bar to be stepped over, not a high bar to be vaulted. Calling attention to this fact and shifting our attention and the discussion from mere ‘compliance’ with the letter of the law to security, as measured by real-world assessments and results, might encourage organizations to consider their needs, obligations and priorities above and beyond regulatory compliance.

3.6 FIXING LAWS AND POLICY

What might the future look like? At a high level, we think that governments, law enforcement and the IT security industry need to take a cue from the public health sector and move (quickly) toward tangible and actionable recommendations that are based on a factual analysis of the cybercrime threat. The US government, with its huge IT infrastructure and resources in the FBI, FTC, intelligence, military and civilian agencies, is in an excellent position not just to compel ISVs to improve the security of their products, but also to serve as an impartial provider of security and threat-related data that IT professionals in the private sector can use to inform their own decisions, free of vendor spin. Too often, however, the government has been a laggard rather than a leader – bogged down by byzantine contracting rules and beholden to a short list of major contractors and systems integrators to provide much of the heavy lifting. Still, there is evidence of innovation and thought leadership, especially where the government has found ways to harness private-sector innovation to serve its own needs and interests. The CIA-backed VC fund In-Q-Tel is a great example of this.

At the policy level, a stronger and wider-reaching legal framework needs to be established to pursue cybercrime cases across national borders. In the West, the Council of Europe Convention on Cybercrime has been in force for more than half a decade, but

similar treaties linking the West with other hotspots of cybercriminal activity in Eastern Europe, Russia and Asia would go a long way toward denying cybercriminals a safe haven and speeding investigations that cross international borders, as nearly all cyber-crime investigations do. We see the need for tougher regulations to hold carriers and ISPs accountable for identifying and blocking malicious traffic on their networks and stop turning a blind eye to suspicious activity.

The huge reduction in spam and, thus, virus volumes that followed the takedown of rogue ISPs like McColo in 2008⁴⁷ and 3FN in 2009 were proof that coordinated government, law enforcement and private-sector actions at the infrastructure level can do more to improve the health of the Internet ecosystem than a million spam gateways deployed in a million customer DMZs. Efforts to formalize such relationships and expand their purview internationally and to conduct regular audits of ISPs wouldn't solve the problems of 'bulletproof' hosting and the offshoring of threats, but it would make it less likely that threats and attacks would emanate from the US, EU and other industrialized nations.

Cross-industry and private-public partnerships and collaboratives also need to be rationalized and made more holistic. We count almost half a dozen such groups – the Anti Phishing Working Group, the Anti Spyware Coalition, Digital PhishNet, MAAWG, etc. – that might easily be phased out or folded into a larger, comprehensive group that simply has 'threat intelligence and response' as its charter.

Rather than heaping on new requirements, industry and government regulators need to get out the red pen and rationalize existing regulations – assessing whether what seemed like must-have technology a decade ago might comfortably be dropped from the list of requirements today. Regulators need to assess whether these technologies are, at a higher level, still adequate to address the problems they were created for, and whether their prescriptions are still relevant in today's marketplace.

End users, meanwhile, need to seek out non-vendor sources of threat intelligence and not merely accept what their vendors or contractors are telling them. Blogs, podcasts, research conferences, consultants and even security intelligence services can make them informed consumers and help deflate the FUD or promises that still drive too much of the IT security conversation.

Finally, vendors need to continue investing in research and development regarding future threats and instill a high standard of intellectual honesty in their messages to customers and the public. As we've noted, the brouhaha over the Conficker worm's April 1, 2009 'ticking time bomb' was great PR for Symantec, McAfee and a few other connected firms, with the story landing on *60 Minutes*, among thousands of other outlets. It was also a rather disingenuous bit of FUD flinging. The absence of any large-scale disruption stemming from the Conficker worm on that day no doubt convinced many that such warnings are often empty. It also overshadowed what might have been

47. <http://www.securityfocus.com/brief/855>

a more measured but equally serious discussion of the challenges that consumers and companies face with the latest generation of threats such as Conficker. An industry-wide effort to check the sensationalism and speak, whenever possible, with a clear voice and from a position grounded in facts would go far in restoring public trust.

SECTION 4

An Enterprise Cybercrime Toolkit

In the previous two sections, we've sketched out – in broad strokes – the dimensions of the cybercrime and APT problem. We've also talked about some of the challenges that enterprises and other organizations face in addressing the changes wrought by organized cybercrime, state-sponsored cyberespionage and the like. We noted in Section 3 that major obstacles include the paucity of reliable data on threats and attacks, poor alignment of compliance mandates and existing enterprise security investments with a new generation of threats, and the muddying effect of industry-generated FUD.

They are all important issues to consider as we survey the enterprise security landscape and try to properly size the problem of cybercrime and state-sponsored espionage. But problems like poor threat visibility or industry-driven FUD are beyond the scope of individual enterprises to tackle. A more pragmatic question for enterprises and for vendors alike, then, is: what changes follow on from the changes in the threat landscape and attacks? If we filter out the FUD and focus on the problem, what technologies or processes do we think will become more (or less) valuable over time? Where should enterprises invest now to stay ahead of threats (or just keep up with them), rather than reacting to attacks and fighting the last war?

We realize, of course, that asking these questions in this way risks oversimplifying the problem. As we've already noted, compliance has eclipsed objective considerations of risk and exposure as a prime motivator for security purchasing. Still, APTs/state-sponsored hacking, malicious insiders and organized cybercrime will be persistent problems in the decade ahead. Organizations of all sizes will have to reconsider their security risks, posture and technology investments in light of real threats, regardless of whether government and industry regulations are well aligned with those problems.

As we spoke with both end-user and vendor experts on the cybercrime problem, one of the most consistent refrains or complaints we heard was that the enterprise response to cybercrime was still far too fragmented. Even in organizations that are sophisticated about fighting fraud and maintaining secure operations and compliance, there might be no clear coordination of efforts against fraud. Security operations might handle threats to the network, and there could be further organizational divisions around the network vs. servers and endpoints. An entirely different group may be responsible for physical security. In large financial institutions, there might be still different groups focused specifically on financial fraud monitoring – anti-money-laundering and so on. Areas like compliance and business continuity might be functionally separate, as well, with different staff, reporting structures, vendors, partners and systems to support those functions. The responsibility for incident response is likely separate, too – involving law enforcement or outsourced entirely to contractors brought in to clean up after a breach.

Organizationally, this siloed approach enables fraud by blinding organizations to multi-vector attacks that combine physical or human targets with software-based attacks or compromises. Furthermore, risk assessments that look only at IP-based attacks or threats from outside tend to underestimate an organization's exposure to risk. Firewalls and IDS are of little use, after all, when attackers can simply tailgate into your headquarters or a branch office and walk out with an employee's laptop. So, at the very least, enterprise efforts around fraud protection need to be informed by a holistic assessment of risks, coordinated across business functions and driven by a single individual. Very few of the enterprises we've spoken with are there yet, though we've seen the beginning of efforts to bring together functional groups with a hand in deterring fraud. Assuming that enterprises can tackle those organizational challenges, they will need platforms to support them. These include monitoring, detection and management tools that help facilitate a cross-functional and enterprise-wide view of fraud.

What's needed, then, is a new enterprise architecture that addresses both legacy security issues and compliance, as well as the danger posed by sophisticated, organized cybercrime, APTs and similar phenomena. This architecture will be built on existing investments and will leverage existing control points, of course. But it will require changes as static or ineffective technologies are phased out and replaced with those better aligned with threats. What might such a platform look like? At a high level, we see the need for the following changes:

- **A hardened enterprise endpoint** – Despite the clear and evident de-perimeterization of the enterprise and the increasing importance of endpoint security, protection of enterprise endpoints and servers is a sore point for most organizations, rooted in an aging and increasingly irrelevant threat paradigm. To change this, we believe the already well-established trend toward convergence of endpoint management and endpoint security needs to accelerate. Enterprises need better visibility, more granular controls and a more modular platform upon which to build protections, as well as more tools for securing data and transactions in environments where compromises are assumed to have occurred.
- **Improved enterprise threat intelligence and correlation** – The value of generic protection against commodity threats may have been adequate in the past. Today, however, the advent of APTs, targeted attacks and sophisticated cybercrime makes the notion even of a 'patient zero' unacceptable – because the cost of compromise for that single customer is likely to be intolerably high. What's needed is improved intelligence on new, emerging and targeted attacks that's tailored to individual customers. Beyond that, enterprises need more actionable reputation information and the ability to correlate both threat intelligence and reputation with their policy and protection tools. Finally, they need threat intelligence services and tools to correlate that information with other enterprise data feeds, ferret out useful intelligence and then present that information in a way that is actionable.

- **Improved network monitoring, analysis and incident response** – Enterprises are still woefully ignorant of the nature of the traffic that is traversing their borders and moving within their networks. Attackers exploit this myopia both to sustain attacks and to exfiltrate sensitive information from corporate networks under the guise of legitimate communications. Fighting APTs and sophisticated cybercrime requires much closer monitoring of traffic flows from both physical and virtual networks, and the ability to drill into that information – replaying events, correlating suspicious activity between different components of attacks and visualizing the ripple effects of attacks within network environments. When attacks are uncovered, enterprises will need to know more than just what family of malware hit them. They need actionable intelligence on the origin and composition of that attack, the components and attack paths used, the parties responsible and, if possible, their motive and MO. This will necessitate an expansion of the incident-response capability within many organizations.
- **Improved data protection and rights management** – As we’ve noted, in the last decade, bad guys have moved the battleground decisively from owning ‘networks’ to owning enterprise data. The recent attacks against Google, Adobe, Intel and other firms only underscore that very evident trend. Enterprises need to respond by adopting more granular, risk-based protections that are focused on sensitive data and the flows of sensitive information in and out of their network, rather than broad-brush protections of assets that may or may not contain sensitive information. Adoption of information rights management tools, which blend elements of DLP, data encryption, identity management and policy management, is a critical step toward hardening enterprises and enterprise data against APTs and other sophisticated threats.

4.1 HARDENING THE ENTERPRISE ENDPOINT

Endpoint security is perhaps the most urgent area of need in the shifting battle to contain organized cybercrime and APTs. As we noted in our prior discussion, enterprise endpoints and users are the battleground of choice for sophisticated, for-profit cybercrime groups and state-sponsored hackers. The reasons for this are simple enough. As the recent Aurora attacks show, enterprise users, even at highly technical and security-conscious organizations, are still easy marks for both commodity malware and targeted attacks that use social engineering. Beyond that, existing endpoint protections are poorly aligned to address APTs and other sophisticated threats – they’re too oriented to commodity threats and mostly blind to targeted, one-off attacks and malware.

All of this augers for a ground-up rethink of enterprise endpoint security and a shift away from the existing model, which still relies heavily on detection of known threats (that is, signature-based detection) through traditional channels such as email. The enterprises we spoke with complained about the complexity and middling protections

offered by the incumbents in the anti-malware space. Some had gone as far as evaluating alternatives, but none had gone so far as to supplant legacy anti-malware suites and signature-based threat protection as their primary line of defense against malware, including APTs. As we noted in Section 3, some of that reluctance is attributable to compliance mandates (notably PCI). But the security industry also does a good job directing questions about the efficacy of their products in the face of new threats back to a standard set of answers – ‘faster’ signature updates and ‘in the cloud’ threat intelligence. Any talk of wholly new approaches to securing desktops, servers, mobile devices and so on tends to get shut down or passed off as unrealistic.⁴⁸

Despite that, we think that consensus is building around the notion of reforming endpoint security. What that will be isn’t yet clear, but here are some ideas that we see gaining currency among the enterprise end users and thought leaders that we’ve spoken with.

4.1.1 PC LIFECYCLE MANAGEMENT SUBSUMES THREAT PROTECTION

As we noted in our 2010 preview of the enterprise security marketplace for the 451 Market Insight Service, the emergence of APTs and targeted threats is accelerating the convergence of endpoint security products with what’s often termed ‘PC lifecycle management’ (PCLM) – patch and configuration management, application control and so on. We’ve written about a number of recent acquisitions and tie-ups that are data points on this trend, including Symantec’s acquisition of Altiris and McAfee’s purchases of Citadel and Preventsys. There was also Trend Micro’s partnership with BigFix, Shavlik Technologies’ partnership with anti-malware vendor Sunbelt Software and Lumension Security’s partnership with Norman ASA.

The logic behind these pairings is impeccable: patch and configuration management vendors do a great job with proactive protections compliance, but offer little insight into threats and malicious code. Endpoint security vendors have, historically, excelled at identifying and removing threats, but have struggled to offer more than signature updates to prevent exploitation of zero-day vulnerabilities by malware. Today, APTs and sophisticated cybercriminal and state-sponsored attacks have shone a light on the limitations of signature detection. They’ve also increased the need for ‘situational awareness’ of threats and attacks beyond what endpoint-based, threat-focused anti-malware products can provide.

Integration of endpoint security with PCLM capabilities like patch and configuration management and monitoring bridges the arbitrary organizational and territorial boundaries that often separate ‘security’ from network ‘operations,’ while providing a central policy platform for managing the posture of deployed assets, assessing the impact of threats to the network and infections, and pushing out changes that can stem or avert attacks. At the same time, trends like client and server virtualization and mobility inflate the number and diversity of enterprise endpoints that need to be managed, putting even more emphasis on platform support, centralized policy management and configuration controls over threat detection and blocking.

48. We note McAfee’s analysis of Operation Aurora – the coordinated attack on Google, Adobe and other high tech firms. When describing how enterprises can protect themselves, McAfee recommends cranking up your Web browser security settings, updating your anti-malware signatures, doing a system scan and enabling its Artemis file reputation technology. With the exception of the tweaked browser settings, it’s not clear how any of the other recommendations would have prevented an infection by the APT/malware used in Operation Aurora.

As we noted in our analysis of *Trend Micro's partnership* with patch and configuration management vendor Big Fix, the convergence of PCLM and endpoint protection provides enterprises with what they've sorely been lacking: better visibility of their deployed assets, fine-grained management of endpoints, and a robust, scalable and integrated management platform with which to manage an array of functions – both security-related and non-security-related. Signature updates, within this converged world, are merely another kind of change that can be pushed out to endpoints, rather than a separately managed process. For BigFix, the addition of comprehensive threat intelligence, leveraged across a global installed base, provides players like BigFix with an inside track on emerging threats, and potentially new tools with which to counter them.

We already see this dynamic playing out in the strategies of players like Symantec and McAfee, which are attempting to build out a managed endpoint story based on strong policy and configuration management platforms (McAfee's ePolicy Orchestrator and Symantec's Altiris). Vendors like Lumension are also moving from patch management to a bigger PCLM story that wraps features like network discovery with endpoint protection (both anti-malware and application control), patch and configuration management, vulnerability management and so on. Unified endpoint agents, a single policy and management framework, and a plug-gable, services-oriented architecture are all part of the vision – designed to simplify the vendor landscape by pulling more, related functions under a single roof. But integrated security and PCLM also gives enterprises better visibility and more flexibility in responding to sophisticated threats.

4.1.2 APPLICATION CONTROL AND BEYOND

We also look to the offerings of startups like Triumfant, an endpoint-focused vendor that has blended centralized configuration management with signature-less endpoint health monitoring. Triumfant's Resolution Manager tracks standard configurations and applications, grouping them like assets on a network. The technology creates a de facto application whitelist, albeit one that is particular to the endpoint and customer, rather than simply generic. It can then spot and fix unauthorized changes to those systems, including alterations to registry settings, memory tables, files and data. Importantly, Triumfant also offers more sophisticated incident-response capabilities than is typical from endpoint security or configuration management vendors – enabling administrators to trace the changes that follow a compromise and apply 'hot patches' from unaffected hosts elsewhere on the network.

While less concerned with health monitoring and threat detection, Persysent Software has built a good story around provisioning, business continuity and compliance with the ability to restore endpoints to a known good state following compromise or crashes. Finally, firms like Tripwire and nCircle have been able to grow rapidly, even in a down economy, by tapping into demand for security- and compliance-focused configuration and change management for physical and virtual assets. The Tripwire Enterprise product blends application control with configuration and file-integrity monitoring, allowing enterprises to detect modifications to

monitored endpoints, and enabling operations staff to detect those changes and take steps to remediate the unauthorized changes through integration with platforms like Microsoft Systems Center Operations Manager. NCircle's Suite360 has similarly evolved into a platform for both security and operations, aggregating asset, configuration, vulnerability and compliance data behind a single pane of glass accessible to both executives and operations teams.

We also note the still-strong ranks of pure-play application control and application whitelisting vendors such as Bit9, CoreTrace, Signacert and Savant Protection. Their story about proactive blocking of unauthorized code has certainly become far more relevant in the wake of Titan Rain, GhostNet, Aurora and other APT attacks. The ability of application control to block designer malware without a signature, and without relying on fuzzy behavioral-based detection, is part of that. But application control also offers a ready answer to pressing problems like how to secure nontraditional endpoints such as SCADA devices, point-of-sale terminals, embedded devices and non-Windows servers and endpoints – all important both for compliance and to thwart broad-spectrum attacks by APTs and advanced cybercriminal organizations. Although limited to verticals such as retail and financial services in recent years, application control vendors we spoke with say they're addressing a much broader cast of enterprise characters these days – firms in energy, IT, manufacturing and pharmaceuticals.

We noted in our discussion of McAfee's acquisition of whitelisting vendor Solidcore that the stock of application control firms is certainly rising in a post-Aurora environment, and we think that 2010 could see some consolidation in the application control space. Still, concerns about false positives and inflexibility still dog application whitelisting vendors, which walk a delicate line in making their products flexible enough to adapt to the heterogeneous application environments of most enterprise networks, without opening the floodgates to all manner of malicious, grayware or potentially unwanted programs.

As the gap between the capabilities of APTs and the detection ability of endpoint anti-malware agents continues to grow, we expect to see enterprise-focused security vendors put greater emphasis on PCLM and application control features to complement their strengths in threat intelligence and malware detection. The most successful of these attempts will be those that reduce the attack surface of enterprise endpoints and close off the avenue of attack by zero-day exploits. Operationally, this opens the patch 'window' a bit wider, allowing overtaxed IT staff to actually test and deploy critical vendor patches without fear of compromise.

4.1.3 A NEW ENDPOINT PARADIGM: 'PRESUMED OWNED'

Our talks with end users and IT firms serving many of the verticals hardest hit by APTs and sophisticated cybercriminal attacks suggest that another clear need being driven by APTs and advanced cybercrime is a way to ensure the security of data and sensitive transactions on endpoints that are 'presumed owed' – in other words, where the presumption is that malware, of some sort, is already resident. We think that tools for managing such environments are increasingly important elements in the enterprise crime-fighting toolkit.

In just one example, a senior manager in charge of cybercrime prevention at a global banking and financial services firm talked to us, in frank terms, about the challenge his organization faced from customer machines that were known to be infected by malware. Looking for specific viruses and Trojans was a fool's errand, this executive complained, and banks lack the support resources or the authority to push out and manage the security of their customers' computers. In the meantime, screen-scraping banking Trojans targeting the company's customers had evolved to the point of being able to manipulate the data on the websites displayed to the customer in real time, making it nearly impossible to detect that their account balances had been modified by illegal money transfers initiated by the malware.⁴⁹

We note a number of vendors that are already offering products that address this 'presumed owned' scenario, and we are looking for others to follow suit as the gaps between detection and threats become even more apparent. We've written about firms like Israel-based Trusteer that are targeting online banking and e-commerce fraud with a combination of a kernel-mode agent and hosted intelligence service that focuses narrowly on securing the data sent back and forth in high-value Web sessions and transactions. Rather than tackling the thorny problem of rooting out all malware that might be resident on a customer machine, Trusteer focuses just on securing sensitive data and the key elements needed to conduct transactions – using a browser helper object (BHO) to lock down Web browsers (IE, Firefox, Safari, Chrome) in order to secure banking sessions. Behind the scenes, Trusteer also secures key functions in the OS that do SSL encryption and decryption of Web sessions from keystroke to delivery to the customer site, and encrypts cached Web sessions that are stored locally. Because many of its customers face APTs and targeted attacks, Trusteer aggregates attack data from across its customer base as an intelligence feed that can be used to identify the specific threats that target the customer base.

UK-based Prevx is another vendor offering security for contexts – such as online banking and e-commerce – where endpoint integrity can't be assured. The company has roots in the host IPS market, but has built out a behavioral anti-malware story for consumer, enterprise and online banking in recent years. Prevx's secret sauce is aggregated threat intelligence and a signature-less architecture in which agents resident on endpoints profile installed applications according to a number of unique identifiers. That information, as well as behavioral data, is then aggregated in the cloud across Prevx's large installed base of free and premium customers (the company claims 10 million deployed agents) to spot anomalies that indicate malicious behavior. Prevx's SafeOnline product is a browser-agnostic offering that secures sensitive transactions by locking out unknown apps and blocking non-core processes on Windows hosts that might interfere with secure sessions.

49. A write-up in *The Tech Herald* details this behavior in the Zeus family of malware as well. <http://www.thetechherald.com/article.php/200938/4459/Zeus-Trojan-moving-past-anti-Virus-protections>

Other vendors in this space include Authentium, which largely targets consumers with a combination of secure DNS and desktop/browser lockdown. We also note vendors like Blue Ridge Networks, which sells a wide range of products for secure remote access, mainly into the government space. Blue Ridge's AppGuard blends application control, port and device control, and data protection with a focus on protecting enterprise and mobile endpoints from zero-day threats and APTs. Vendors like BeCrypt and RedCannon Security are offering a mix of DLP and secure connect capabilities – both promising bootable, clean environments for secure remote access on endpoints that are 'presumed owned' using USB devices.

Today, these kinds of transaction-focused security agents have clear applications in verticals such as banking, retail and government – but fewer applications for enterprises. We think the relevance of such technology will grow along with adoption of hosted enterprise applications and cloud computing within the enterprise. Indeed, as the trends toward enterprise cloud adoption and enterprise mobility both pick up steam, we expect that demand for technologies such as risk-based authentication and transaction-focused protection such as that offered by Trusteer and others will increase as well. Married to identity management, user profiling and fraud analytics of the kind that firms like Guardian Analytics are selling into the banking space, transaction-based protections like Trusteer could form the foundation of a re-envisioned endpoint protection paradigm encompassing fine-grained user and device profiling and strong authentication to protect users both on premises and in the cloud.

4.2 IMPROVED THREAT INTELLIGENCE AND CORRELATION

If the 'threat intelligence' model, circa 2001, was about finding viruses first and getting your signature out, the picture in 2010 is much more complicated. Knowledge of new and emerging malware is, of course, a necessity. To that end, top-tier security firms like Symantec, McAfee, Trend Micro, Microsoft, IBM and others have long offered threat intelligence feeds that leverage the output of their research labs, as well as global deployments of vulnerability scanners, endpoint and gateway security products, honeypots and so forth. Symantec's DeepSight service is a good example here, while McAfee has its Threat Intelligence Service. Typically, these services combine data on vulnerabilities and new threats that target them, and they are consumed through information portals or pushed out to customers in some XML format file – for use with internal risk management platforms. Different vendors offer different degrees of granularity around these kinds of services, with feeds focused on particular verticals, geographies or even individual companies.

But APTs and sophisticated cybercrime operations have exposed the need for more actionable intelligence. As we've discussed, at the tactical level, modern threats target flaws in both application logic and business logic, not to mention open source intelligence on a target organization, its employees and infrastructure. To respond to such

threats, organizations need a wider range of intelligence services that encompass vulnerabilities within their IT infrastructure, but also intelligence on external threats – sources of malicious activity online, chatter that might be the early stages of attack planning, or information that might impact an organization’s reputation or the safety (or integrity) of its employees.

Beyond that, we believe that enterprises need a way to consume such data and translate it into specific policy actions or remediation. Security portals and Internet-wide ‘threatcon’ ratings might have been useful in the days of global worm and virus outbreaks like Blaster and SQL Slammer, but they’re of little value in an age of APTs and targeted, low-profile attacks. At the same time, log management and security information management platforms do filter out the noise created by disparate security products, but still focus mainly on security ‘events’ at the expense of intelligence. In this section, we take a look at emerging sources of threat intelligence that we think will prove useful in fighting APTs and advanced cybercrime. We also look at some tools that we think enterprises need to invest in to make the best use of the security data and intelligence they collect.

What are we referring to by ‘threat intelligence?’ It’s a broad term that encompasses a wide range of disciplines. At the low end, threat intelligence comprises online reputation monitoring of the type that large messaging and endpoint security incumbents have long carried out. Offerings like Cisco’s (formerly Ironport’s) Senderbase and McAfee’s (formerly Secure Computing’s) TrustedSource reputation services track the sources of malicious and spam email traffic as well as phishing attacks, and use that data to develop online reputations for particular Internet addresses or domains. Today, that kind of information is being combined with a broader range of threat information from endpoints and gateways, as well as globally distributed IDS sensors and honeypots.

Following in the footsteps of other fledgling industries that have sprouted in response to new threats (think AV, anti-spyware and anti-spam), botnet-detection firms like Damballa and FireEye came to market with IP around tracking active bot networks, then spotting stealthy communications between bot-infected hosts on enterprise networks and botnet command-and-control servers hosted outside the firewall. Today, with the advent of fast-flux botnets and other evasion techniques, botnet detection has evolved from merely tracking command-and-control servers to automated detection using machine-learning algorithms to spot telltale signs of botnet infection and generically identify botnets.

Recent years have also seen the emergence of a number of brand- and reputation-monitoring firms that target this problem. Initially focused on phishing prevention via phishing-site monitoring and brand protection, firms like Cyveillance (*now part of UK-based QinetiQ*) and MarkMonitor have evolved into diversified businesses whose value proposition lies in their ability to monitor the amorphous cybercriminal underground for chatter that might represent the early stages of a planned attack, or for detritus (stolen files, customer information) that points back to an undetected compromise. Cyveillance, just to take the most prominent of these firms as an example, today offers a range

of services: licensing threat intelligence feeds to a wide range of security vendors (Blue Coat Systems most recently) that filter against the URLs linked to malicious content, operating a research lab that does malware analysis, and offering phishing-site takedown services, as well as more specialized brand monitoring (misuse of trademark, etc.), distribution services and threat surveillance (such as executive protection services). Since being absorbed into QinetiQ, those capabilities are being integrated into a managed security offering targeted at the government sector.

We also note smaller outfits like Team Cymru and Cassandra Security, which have found a niche by leveraging both technical expertise and access to provide premium threat-intelligence feeds. For Cymru, those feeds are around malware, phishing and botnet command-and-control networks. Cymru partners with ISPs to analyze large volumes of threat and attack data and with Interpol and law enforcement in 62 countries on malware and phishing-site takedowns. The company's associates infiltrate and monitor cybercrime groups and information-sharing hubs and do reputation monitoring for select customers to spot nascent attacks or merely chatter that might amount to the early stages of planned attacks. Cassandra's experts are focused more on critical infrastructure as well as cyberwarfare and cyberterrorism.

The problem here continues to be finding an audience for specialized fraud intelligence, as well as niche protections like botnet-detection services, beyond the large enterprises, government, financial services firms and ISPs that are the current audience for such services. Concern about cybercrime, state-sponsored hacking and APTs could provide more justification for a discrete investment in threat intelligence – as enterprises look to get a handle on what malicious code has slipped past their gateway and endpoint monitoring tools. The other problem is that such services aren't mandated by state, federal or industry regulations, making them discretionary investments that must take a back seat to other, mandated technologies.

A more likely scenario is for threat intelligence and reputation monitoring to be absorbed into larger suites of products and managed or professional services. Qinet-iQ's acquisition of Cyveillance is just one data point in that trend. We note other, larger vendors that are also wrapping fraud and threat intelligence into relevant products and services. Security and storage giant EMC/RSA's FraudAction Intelligence service is being offered to select customers and provides information on targeted exploits and threats, as well as breach notification. RSA says it has also expanded its threat research expertise around Trojan analysis and offers phishing-site/attack-site identification and takedown services as well.

After spinning off much of its managed security portfolio to managed security services provider (MSSP) SecureWorks, VeriSign is leveraging its role as a critical-infrastructure provider (it hosts the root servers for the '.com' domain) and focusing its efforts on closely related areas such as DDOS protection where firms like Prolexic and ATT play. The company has also grown its anti-fraud offerings with a variety of services. Today, VeriSign offers both protections for online banking and AML, as well as a variety of

incident management and response services, including malicious-code analysis and reverse-engineering – typically delivered in the wake of a breach or sold as part of larger packages of services. VeriSign has also branched out from vulnerability research into reputation monitoring, although it says it offers this as a complement to existing services, rather than competing head-to-head with Cyveillance, MarkMonitor and other firms.

We also see the potential for botnet detection to become one element of a larger anti-fraud and cybercrime suite that also encompasses DLP, reputation monitoring and other services. For now, firms like Damballa and FireEye have broadened their reach (or at least their marketing message) to encompass APT and zero-day threat detection and analysis as well as data-theft detection and incident response. That message could resonate with larger vendors in need of domain expertise around botnets, command-and-control networks and APTs.

4.2.1 BROADER APPLICATIONS FOR ANTI-FRAUD

As we went out and talked to end users at organizations that are facing serious and sophisticated hacks and cyberattacks, it was often issues around the management of identity that came to the fore. Indeed, in the age of insider threats, sophisticated social engineering and APTs, a better understanding of the identities (users, employees, customers, business partners) that interact with your organization is a prerequisite to stopping many flavors of cybercrime, or spotting slow-moving, multistage attacks by APTs. CSOs told us that they need more clarity about the identities at work within their environment, not just their IT infrastructure. Increasingly, the context in which employees and customers interact with that infrastructure is important to spotting and thwarting sophisticated attacks, account takeovers and so on. The goal here is to make sense of the ‘single user, multiple identities/roles’ mess, and put security events in the context of specific identities, to spot both compliance violations and nascent threats. Spotting sophisticated hacks and cybercriminal activity is about correlating information on users, data on new and emerging threats, external reputation (IP-based and otherwise), fraud intelligence and so on.

But the inability of current identity management systems to make nuanced access decisions based on policy and business context is a major shortcoming, and one that bears directly on the cybercrime and APT problems. Just as an example, the CSO at the electronics payments provider we interviewed talked about how his organization had correlated a list of phishing-site IP addresses to login attempts for his company’s Web-based payments platform. Of the 100,000 or so suspect login attempts that analysis produced, the company identified 1,000 successful logins and began close monitoring of those accounts to identify fraudulent patterns. It also introduced additional checks, such as voice confirmations, to prevent unauthorized money transfers. The problem is that this was mostly an ad hoc and customer-initiated process conducted after the fact.

A better solution would be to evaluate access decisions in real time along with externalities like whether the location of the access request matches up against a suspect IP address, and then modify the access decision based on those factors. To do that, enterprises need to be able

to understand patterns of user activity and access to data, and discern how different user identities diverge from ‘normative’ behavior in ways that connect them with each other as part of a sophisticated fraud ring or APT-based attack. We think of this as a kind of CRM for fraud (fraud relationship management, anyone?) that could spot suspicious trends and correlate multiple ‘malignant’ identities as part of a single threat linked to an individual or group.

While ‘FRM’ might not be an acronym we see thrown around anytime soon, we note that there is a precedent, broadly speaking, for marrying identity, activity monitoring, reputation and threat intelligence in ways that might enable enterprises in the fight against sophisticated cybercrime and APTs. If we look to the banking and financial services verticals, we see a range of fraud-monitoring tools to flag suspicious behavior that may indicate fraudulent activities such as unauthorized money transfers or activity associated with money laundering, like ‘structuring’ – arranging a series of payments in a deliberate attempt to avoid transfer limits. The commonality among these disparate technologies is sophisticated business logic and artificial intelligence that can correlate structured and unstructured data from core transaction-processing systems, analyzing activity from brokerage or check processing, ATMs, automated clearing houses and wire transfers. Other layers can then be laid on top of that, including log and event data, customer data, geographic information, global threat intelligence and so on.

While this is a nascent trend, we think there’s crossover potential here, as organizations outside of banking, financial services and e-commerce look for tools to help spot sophisticated patterns of fraud, espionage and other malicious activities. Juniper Networks is one company *that has joined* threat management to identity management offerings across a number of product categories (SSL VPN appliances, SRX gateways, Network and Security Manager). Juniper’s appeal is to allow enterprises to consider user role, along with endpoint state for pre-admission network-access decisions (fka ‘NAC’) and post-admission resource-access requests.

We’ve also talked about the growing focus of vendors like IBM (through its Tivoli group) and CA (*with its Eurekaify acquisition*) on a broader problem of identity governance and ‘continuous risk management’ – which involves both identity correlation and the job of putting identity in context with transaction data, security policy and compliance objectives. We note firms like Guardian Analytics that are targeting problems like account takeover – mining banking application and session data and running it through a proprietary risk engine to create a profile of individual users that can be used to spot suspicious activity. Although focused on online banking as a use case, Guardian Analytics’ FraudMAP has applications outside of the banking vertical, including e-commerce or other applications.

Farther out on the periphery, we see enterprise applications for technology from startups like Memento Security that sell next-generation tools to spot insider threats to the banking industry around problems like deposit account fraud – a category that encom-

passes things like check and online banking fraud. Memento's claim to fame is AI that can make sense (using probabilistic models, Bayesian belief networks, etc.) of complicated environments with mixes of known and unknown factors. That's useful in spotting 'collusive networks' of individuals and businesses that, the company acknowledges, could have applications beyond financial fraud.

4.3 LEVERAGING NETWORK MONITORING AND ANALYSIS

Attacks like Titan Rain, GhostNet and Aurora have also underscored the degree to which APTs challenge the traditional layered approach to network security, comprising perimeter, network and host protections. Modern threats leverage on-premises malware and activity along with off-premises, Internet and P2P-based command-and-control components, drop-sites and other staging resources. They're also adept at exploiting the blind spots in existing security deployments, communicating over common ports and protocols that pass through firewalls and receive little scrutiny.

Spotting such threats requires more sensors at the network edge and at the network segment, as well as on servers and endpoints. Enterprises need a more sophisticated understanding of network traffic flows across and within enterprise borders, as well as more powerful packet-analysis tools to spot malicious payloads and sensitive data that is traversing their networks and might be indicative of malicious activity.

4.3.1 GROWING EYES AND EARS: NETWORK TRAFFIC ANALYSIS, DLP AND BEYOND

Today, a wide range of vendors offer variations of this kind of capability. A long list of IDS and IPS vendors such as IBM/ISS and McAfee do basic protocol analysis as part of their threat-detection offerings. Beyond that, network traffic and protocol analyzers are offered by firms such as Niksun, NetScout, WildPackets, Fluke Networks and Wireshark, which have historically focused on network performance issues and fault testing.

We've noted in recent years the emergence of a newer generation of players that are applying the same underlying technologies but adding more in-depth network security analysis. The growing problem of APTs, sophisticated cybercrime, state-sponsored hacking and online fraud have only heightened demand for the kinds of sophisticated capabilities for network traffic capture and packet analysis that these products offer. One firm that's targeting the APT problem with sophisticated full-packet-capture capabilities is NetWitness, which landed a funding round from Summit Partners in January 2010. The company grew out of a project for the US intelligence community, and says that its business is expanding rapidly with interest not just from the federal government sector, but also from critical-infrastructure providers, financial services firms and others.

NetWitness promises enterprise-wide real-time, full-packet capture, storage and analysis. Its products use patent-pending analytics capabilities to de-obfuscate advanced threats by allowing analysts to decode network traffic and understand, frame by frame, how Trojans, rootkits, Web-based attacks and other threats are working. For example, NetWitness' Capture and Decoder products can store and disassemble Web pages that might have been used in drive-by download attacks, allowing analysts to replay Web sessions frame by frame, isolate page elements that might contain exploits, and then run those exploits in a sandbox environment to study their behavior. Integration with AD and other user stores also allows NetWitness to view activity associated with a specific user or IP address, as well as trace threats back to their origin using geolocation. Links out to security information and event management (SIEM) platforms allow the users of those tools to tap into the deeper analysis – for example, viewing full HTTP sessions or identifying a malicious command shell masked in DNS traffic.

Similarly, Solera Networks leverages a proprietary file system and packet-capture and storage technology to do ultrafast (10Gbps) and comprehensive network capture, indexing and analysis. Although lacking the sophistication of firms like Niksun or NetWitness in decoding and analysis, Solera has forged partnerships with vendors like Palo Alto Networks, Sourcefire and ArcSight – matching up alert information with detailed, captured session data associated with it.

Network-based DLP from firms like Symantec/Vontu and Fidelis have a role to play here as well. Fidelis' XPS product, to take just one example, uses sensors both at the gateway and inside the network to discover and classify in-network traffic across all ports and protocols, while also monitoring sensitive channels like mail/SMTP and HTTP. Unlike NetWitness, Fidelis works in-line, offering full session awareness and leveraging proprietary data decoders that strip off application wrappers to do data fingerprinting and keyword analysis. Multidimensional content profiles can then be linked to specific policies governing data movement. For example, traffic can be shut down, or files can be quarantined if Fidelis' scanners, packet sniffers, TCP session trackers, etc., catch something amiss.

Picking over the details of attacks like Aurora, in which sensitive intellectual property and source code was siphoned from high-value servers and endpoints, some combination of DLP and network traffic capture would seem to be in order, and we note that firms like Solera say they are looking to relationships with DLP vendors to provide deeper context and analysis for their traffic capture and indexing capabilities.

4.3.2 DRILLING DOWN ON THREATS: IMPROVING INCIDENT RESPONSE

In terms of endpoints and servers, APTs and sophisticated cybercriminal attacks have increased the need for more in-depth incident-response tools including forensics and root-cause analysis as well, as organizations hit by such attacks try to understand more about their attackers and the tools they used, so they can better understand the ramifications of the attack for their networks and data.

We saw elements of this dynamic play out in the aftermath of the Aurora attacks, as details about the attack emerged from those companies brought in to assist with forensic analysis of it – McAfee, VeriSign/iDefense and so on. We also saw close scrutiny of the components of the malware from experts like Joe Stewart at SecureWorks to attempt to link the attack to a group or nation-state.⁵⁰ Although merely the public leaks of much more in-depth analysis of Hydraq and other APTs and malware by a range of players (government and otherwise), the focus on the provenance of the Aurora threat is indicative of a broader trend. Enterprises that have been targeted by malware don't merely want it removed, they want to understand its workings and origins. Beyond that, security firms are looking for ways to work backward from post-infection malware forensics and the data gleaned from incident response to proactively inform policies and rules that can block future attempts to compromise networks and valuable assets. Today, we see a number of firms offering deeper, full-spectrum threat intelligence both in the context of professional services engagements (mostly), but also as a consumable service to be leveraged by vendors further up the security intelligence food chain. EMC/RSA tells us that it has ramped up its incident-response and malware forensics capability, with particular emphasis on analysis of Trojan horse programs, as part of a multi-front effort to increase its fraud intelligence offerings and integrate those more closely with its other security products and services.

Mandiant is another fast-growing firm that works extensively with government, financial services and other targeted verticals. The company, which sells incident-response software services, has been a thought leader on the APT problem and is beginning to leverage partnerships with third-party firms (we note a deal to license Bit9's global software registry) to hone and speed up its incident-response capabilities.

HBGary is another leader in this space. The firm has leveled proprietary memory forensics technology developed while researching rootkits to create a suite of incident-response and forensics tools. By understanding key elements of Microsoft Windows memory management architecture, HBG has developed tools that can reconstruct captured Windows images (including VMs) with total accuracy and then step through program execution at a granular level, showing memory allocation, library and processor access,

50. Stewart claimed in a SecureWorks blog post that a cyclical redundancy check routine used by Hydraq was unique and of Chinese origin, providing forensic evidence that strengthened Google's claim that the attack originated from that country. Subsequent information, however, called into question whether the routine Stewart identified really was as rare as he claimed, or unique to the Chinese software development community. See: http://www.theregister.co.uk/2010/01/26/aurora_attack_origins/

etc. – and then use that unique information to fingerprint malware executables, changes linked to malware infection or other activity, and extract forensic information from memory post-infection. HBG's main appeal so far has been within the government sector and among large financial institutions, where interest is high in understanding the provenance of threats, as well as their inner workings. But the company has also leveraged that research to develop Digital DNA, a signature-less method of detecting malware that uses behavioral-based malware identities that it sees as being more commercially applicable.

We think the direction taken by firms like HBG is illustrative of the direction of other forensic analysis and incident-response vendors. The kind of analysis and detection capabilities they offer will also become a more valuable commodity as a wide range of security players zero in on the problem of hard-to-detect persistent malware and APTs.

4.3.3 PULLING IT ALL TOGETHER: SIEM AND EVENT CORRELATION

As we've noted in our research, enterprises are increasingly looking for ways to integrate security event and log data for the purposes of both threat modeling and forensics as they delve into sophisticated cybercriminal attacks and ferret out APTs or malicious insiders.

We've written about *the drive toward convergence* of previously disparate technologies like log management and SIEM. In these scenarios, log data provides the grist for establishing a horizontal view of network assets and events, while SIEM platforms provide correlation, visualization and presentation to help analysts and IT staff make sense of that data and evaluate it in the context of risk and compliance objectives. That trend has been driven by threats and demand for greater automation of IT management functions. The value, today, for enterprises is to help orient their security investment and response around risk, and to reduce the complexity and cost of compliance and IT management. However, as incidents like the Aurora attacks shift attention to problems such as APTs and sophisticated cyber-crime, we expect the relative importance of broader-spectrum threat correlation to grow and to push demand for new tools to help IT staff correlate and spot patterns in structured and unstructured data from endpoints, dedicated security appliances, network taps and so on.

This certainly looks and sounds a lot like today's SIEM platforms, and we think that such products are already moving in the direction of more holistic threat intelligence platforms. As an example, in recent months, vendors like *ArcSight* have added features that begin to span some of the silos we mentioned above: integrating unstructured application and user data with their existing SIEM and log management capabilities, while broadening the capability of those platforms to correlate data from structured log events with sources like identity management systems. Normalizing, storing and indexing the data exposes it meaningfully to queries. (Take, for example, the malicious insider whose workstation is linked to failed login attempts on a sensitive internal application, and who is also slipping sensitive files out over webmail or Skype.) Similarly, EMC/RSA has built out existing investments in SIEM (enVision), adaptive and risk-based authentication technology (Passmark, Cyota, Verid) and its threat research (to combat phishing, pharming and Trojans) into a comprehensive

fraud management platform. We expect to see other SIEM and log management vendors moving in this direction, as well, as concerns about sophisticated hacks, cybercrime and breaches take a seat next to compliance as a driver of new business.

While ArcSight and EMC are out front in this trend, smaller competitors aren't far behind, with plans to bring security and IT operations closer together around the problem of compliance and advanced threats. Vendors like *NitroSecurity* have used acquisitions to branch out into adjoining areas like database transaction monitoring and DLP in search of a broader IT-GRC story that encompasses threat correlation across multiple channels. Companies such as Tenable *have moved upstream* from vulnerability scanning to log management and SIEM, in pursuit of a broader story about a unified IT-GRC platform. *Intellitactics* has focused on building hooks into policy management platforms like McAfee's ePO. TriGeo and eIQnetworks, as well as Splunk and LogLogic, are worth watching here. In some cases, these offerings are being offered as hosted services (as with Alert Logic) or through MSSPs (netForensics, Tenable).

But SIEM and event correlation only help to sort out the noise from existing security investments and other deployed technology. One of the major challenges in investigations of cybercrime and APTs is making sense of the data supplied by SIEM platforms as well as by the kinds of packet-capture and DPI tools from NetWitness and others. Today, we note players that are doing just that: joining up SIEM and IT search (to use Splunk's terminology) with more specialized functions like computer forensic analysis, network topology mapping and threat intelligence in search of better information analysis tools that can spot cybercrime and the work of APTs, and help understand the dimensions of attacks.

We also note the emergence of new tools that are designed specifically to drill into the terabytes or petabytes of network-traffic-capture and threat data that wash up in the course of a modern incident investigation. Palantir recently received press attention for the firm's role in mapping the GhostNet attack on the Tibetan Government in Exile and helping forensic analysts understand sophisticated, state-sponsored attacks on US government infrastructure. The company offers products for both the government and finance verticals that combine structured data (access management, packet sniffers) and unstructured data feeds in a common (Oracle) data store. We also note Lookingglass Cyber Solutions, which offers what it calls 'situational-awareness and incident-response tools' that are geared toward spotting APTs and other sophisticated attacks. The company's technology does network topology mapping that allows analysts to visualize the assets deployed on their network, analyze netflow and other data, and connect suspicious and malicious activity internally to data on external threats such as botnets and phishing sites, geo-IP data and so on.

Clearly, there's overlap between these kinds of attack-mapping and situational-awareness tools and what firms like NetWitness and Solera offer. The difference may be one of degree, with platforms like Palantir's offering a layer of abstraction on top of the infrastructure-level data and netflows collected by a platform like NetWitness. Platforms like

Palantir's provide analysts with near-infinite flexibility to query, parse and slice data in order to discover patterns between different attack elements. As an example, analysts using Palantir can create dynamic links between IP addresses, attacker profiles and other data, such as geolocation, as they investigate incidents and combine the work of multiple analysts working on investigations, model networks and attack hypotheses, and so on. In the long term, they could become the platform upon which the heretofore separate disciplines of fraud analysis, network traffic analysis and SIEM converge – providing something that looks like a holistic enterprise cybercrime-fighting toolkit.

4.4 PROTECTING DATA: A SHIFT TO INFORMATION RIGHTS MANAGEMENT

Closely related to the need for better identity governance is a need for more granular data protection informed by identity and policy. We (and others) have loosely termed this 'information rights management' (IRM), and we see it as an increasingly important tool in the enterprise cybercrime-fighter's toolkit. The logic here is easy enough to grasp: fast-evolving and sophisticated threats and an evaporating network perimeter make environmental protections harder to maintain. That, in turn, pushes the need for protection down to the data level.

That kind of logic has driven investment in DLP technologies that can both fingerprint and secure sensitive data at rest and in motion, or close off loosely monitored egress points such as local USB ports. But DLP has limits. As we've noted here and elsewhere, DLP's strengths in data discovery and classification often don't extend to the policy layer. At the same time, enterprise rights management does a fine job of managing user roles and access privileges, but often lacks granularity at the data level. As we noted in our *2010 M&A Outlook report*, IRM provides a kind of connective tissue for embedded enforcement, bringing together elements of digital rights management, DLP and classification, encryption and key management with an identity and access management (IAM) layer for document-level controls.

How does this fit into the cybercrime-fighting toolkit? We see IRM as a tool for enterprises to move from protecting 'assets' to more granular, risk-based protections that are focused on sensitive data and the flows of sensitive information in and out of their networks. As incidents such as the Aurora attacks remind us, intellectual property, state secrets or sensitive financial/identity information is often the target of sophisticated cybercrime and state-sponsored cyberespionage. Enterprises that don't have visibility into where their sensitive data lives – and tightly controlled policies on how it is secured, who has access to it and where it's moving – risk not only falling short of compliance goals, but becoming victims to APTs and targeted attacks aimed at the very heart of their organizations.

Today, we see vendors like CA blending the DLP and data-discovery capabilities it acquired via Orchestra with archiving and log management, as it focuses on compli-

ance dollars. MSSP Trustwave is building a tightly integrated network and endpoint compliance offering that encompass device and data discovery through a rich, policy-based data-protection offering. *(That company has acquired both DLP with Vericept and data discovery and encryption with BitArmor in the past year.)* Expect larger vendors with investments in endpoint, network infrastructure or IAM to follow suit. Liquid Machines, GigaTrust and Varonis, as well as InDorse Technologies and transparent encryption vendor ZafeSoft, all have pieces of the IRM puzzle. In the DLP space, we note Fidelis Security Systems, Verdasys and Safend, which has added data discovery and classification capabilities.

SECTION 5

M&A Opportunities

Having outlined in broad strokes the enterprise cybercrime-fighting toolkit, the natural next question is what the opportunities for disruption are – areas where nimble startups can gain traction and incumbent vendors may look to strengthen their hands.

5.1 BUILDING A CYBERCRIME-FIGHTING STACK

Looking across the various technologies in Section 4's discussion of the enterprise cybercrime-fighting toolkit, we see a new kind of technology stack emerging in the enterprise security space around the problems posed by sophisticated APTs and state-sponsored and organized cybercriminal operations.

If the existing enterprise security stack builds up from network monitoring and protection through to OS and application security, this cybercrime-fighting stack is built upon better, richer data about threats and malware – not just their behavior, but their very makeup. This includes sources of the attacks as well as the structure of organized cybercriminal groups or state-sponsored actors behind the attacks, and encompasses external and internal threat intelligence, reputation monitoring, and sophisticated network traffic capture and analysis.

The job of building this stack will fall to a variety of large vendors, managed services players, infrastructure providers and systems integrators with an interest in addressing cybercrime, APTs and hacks sponsored by nation-states. The following sections outline some of the opportunities we see for disruption and consolidation as vendors scramble to assemble pieces of the cybercrime-fighting technology stack.

5.2 THREAT INTELLIGENCE FIRMS FIND SUITORS IN THE WINGS

More than a year ago, we *picked 'forward threat intelligence'* as one of a handful of 'recessionistas' – areas of continued investment that would flourish despite a down economy. With the focus of attention more than ever on APTs, we expect to see talent and tools at a premium when it comes to conducting threat intelligence and spotting threats to critical infrastructure, reverse-engineering APTs, fraud detection and incident response.

In the short term, a menagerie of small consulting and products firms that specialize in these areas will likely be absorbed into larger security firms looking for visibility into the criminal underground and the tools and talents for monitoring nascent attacks and scams. We have already seen acquisitions like QinetiQ's purchase of Cyveillance in May 2009. Adjacent firms like MarkMonitor, BrandProtect and Envisional could also be brought into larger, diversified security and infrastructure players, systems integrators or MSSPs that seek tailored security services, or they could complement existing threat research and on-premises offerings.

We wrote in October about *Cisco Systems'* need for better threat intelligence to supplement the SenderBase and SensorBase online reputation technology as it looks to build value across its security portfolio. This is also true for competitors like F5, Fortinet and Barracuda, as those firms look to broaden the capabilities of their security point products and unified threat management (UTM) offerings. MSSPs like SecureWorks and Trustwave, or larger, multinational players like BT and Verizon Business could also scoop up smaller players in an effort to build high-touch security services that blend reputation and brand monitoring with more traditional outsourced offerings.

5.3 MALWARE FORENSICS: EVERYONE'S SECRET SAUCE

While most of us were off counting Blaster worm variants, experts like Greg Hoglund were quietly warning about the dawning problem of rootkits and other kinds of advanced threats impossible to detect by conventional means. Fast-forward five years, and firms like Hoglund's HBGary, Mandiant, Cassandra Security and others are at the forefront of the converging worlds of malware forensics and the more loosely defined 'threat intelligence.'

At a baseline, HBGary specializes in forensic tools that help analysts understand the behavior of malware, post-infection, in minute detail. But the company has begun aggregating its forensic intelligence into more genotypic malware definitions that can be used to detect sophisticated, zero-day threats in the absence of a signature, and to categorize and rate those threats by severity.

As Hoglund points out, there are hundreds of thousands of different keylogger programs, but only a handful of ways to sniff keystrokes on a Windows system. HBGary focuses on the latter, rather than the former, in identifying threats. That kind of capability, in addition to HBGary's deep pool of experts and knowledge base of threats, is an increasingly valuable commodity for diversified vendors that want to boost their professional services capabilities (McAfee?) and also increase their understanding of advanced malware and the cybercrime underground.

We are already seeing activity from vendors in the threat-analysis space. Palantir and HBGary recently announced a partnership to feed HBGary malware research directly to Palantir's console, allowing analysts using that platform to benefit from granular data on malware behavior and origin as part of their full-spectrum threat analysis. We'd also note integrations with McAfee's ePO platform, allowing that company to more accurately score the severity of threats based on their behavior rather than signature-based identification.

While we expect the market for forensic tools to remain small and highly specialized, we think any number of security companies are interested in scooping up both the tools and expertise that firms like Mandiant, Cassandra and HBGary have in the analysis and disassembly of Trojans, rootkits and other next-generation malware. Look for interest not only from incumbent anti-malware vendors, but also from systems integrators with significant government contracts in 2010.

5.4 SIEM TAKES CENTER STAGE

One clear trend is the need for more powerful tools to correlate structured and unstructured threat information across an enterprise IT infrastructure, as well as external reputation and threat intelligence feeds. We've written in the past about the first phases of the shift toward a more integrated IT management and threat intelligence capability, as we've tracked the gradual *convergence of SIEM and log management*. The coming together of those two feature sets is driven by the need for broader monitoring of security infrastructure (log management's traditional value) and better identification and visual representation of threat trends (SIEM's strong suit).

The outcome, as evidenced by category-leading vendors like ArcSight, is a platform that marries heterogeneous event capture, log management, correlation and graphic presentation/visualization. Such tools can give IT security staff and analysts a hierarchical view of threats and vulnerabilities that correlates with risk and compliance objectives. As noted in the previous section, however, more is needed to address the challenge posed by APTs and sophisticated cybercrime. More recently, ArcSight moved to consume unstructured data such as IM and Web content as it targets multivector fraud within enterprises. We expect to see continued broadening of that palate as SIEM vendors look to pull more relevant security and threat data into their consoles.

In the short term, there are a number of interesting trends that could lead to disruption and M&A. At a base level, a slew of vendors are gaining mid-tier and large enterprise traction with offerings that span the security-IT operations gap, combining elements of SIEM with log management, vulnerability scanning and other relevant bits. In addition to ArcSight, we note LogRhythm, Q1Labs, NitroSecurity, SenSage, Intellitactics, TriGeo, Tenable Network Security and eIQnetworks, Splunk and LogLogic. The bigger trend may be large security and infrastructure vendors getting into the SIEM space. As we've noted, Cisco needs more than just a refresh of its SIEM capabilities to replace the CS-MARS product.

Anti-malware vendors McAfee and Symantec have invested heavily in policy and configuration management in recent years. The mature endpoint suites of both vendors would provide endpoint visibility and enforcement – features that the current crop of SIEM and log management vendors lack, but that trends like mobility and virtualization demand. These vendors still have noticeable gaps to fill in the arena of SIEM – an increasingly evident hole as they vie for lucrative government and large-enterprise contracts where sophisticated attacks and APTs are a top concern. A mature vendor with a big federal footprint like ArcSight could be in the sights of both firms (in fact, we've heard rumors of just such a union).

Enterprise-focused anti-malware or security and compliance vendors like Sophos or nCircle could follow suit – or even jump out ahead of this trend, picking off a smaller diversified security vendor with promising log and SIEM capabilities (Tenable? NitroSecurity?), or a close cousin like the ever-popular Splunk, to act as a management layer that could amplify the value of existing vulnerability-scanning and threat-detection products in areas such as cybercrime and fraud prevention as well as compliance.

EMC/RSA's recent acquisition of governance, risk and compliance (GRC) specialist Archer Technologies is yet another wrinkle: the addition of a process management layer on top of the company's enterprise security information management, risk-based authentication and DLP technologies. The idea is to broaden the ecosystem of third-party products and data sources that can be brought to bear on decisions about threats, risk and compliance, providing analysts with tools to prioritize incidents and translate them into specific policy recommendations and remediation actions. CA is following a similar course as it pulls its identity, DLP and log management wares together. Smaller GRC players like Modulo Security, Skybox Security and Agilance could come into play if larger vendors see merit in EMC/RSA's strategy and look to follow suit.

Finally, we see opportunities for successful managed and hosted SIEM offerings, given the continued complexity and cost of deploying, tuning and managing platforms like ArcSight. Firms like Vigilant (which partners with SIEM vendors ArcSight, Q1 Labs, and EMC/RSA) now offer customers a hybrid managed-hosted SIEM. Intellitactics also does managed SIEM partnering with MSSPs like Perimeter eSecurity, Telus and Quest.

On the pure hosted front, there's less to show, although log management vendor Alert Logic has partnered with MSSPs on white-label versions of its hosted log management services, as well as a compliance-focused Log Review Service. Some focused M&A from an acquisition-minded MSSP like Perimeter eSecurity or Trustwave in the SIEM space is a possibility in 2010.

5.5 OPPORTUNITIES FOR (SMARTLY) MANAGED ENDPOINTS

We don't think it's going too far out on a limb to predict a radical change in focus and investment around endpoint security, resulting directly from some of the trends we've talked about in this report. Existing endpoint protection suites are poorly aligned with the current threat landscape: reactive, resource-intensive and difficult to manage.

Subscription-based AV for consumer and enterprise markets is still a popular product, and a lucrative one for large incumbents. But a number of factors are turning signature-based detection into a commodity. We noted as much in our *recent report on 'Freeganism.'* Since that was published, the ranks of anti-malware vendors giving away product has only increased, and shows no sign of abating.

What does this mean? Look for innovation in endpoint security to come from the margins, rather than from established endpoint players. We like the combination of hosted, aggregated threat intelligence with a lightweight agent, behavioral detection and application control that vendors like Trusteer, Prevx and Authentium offer. We expect that model to gain traction with MSSPs and security-conscious verticals that have embraced the reality of 'assumed owned' endpoints.

However, incumbent anti-malware vendors could decide to subsume these upstarts. These smaller vendors (smartly) have always talked about their offerings as complementary to endpoint protection suites, even if it was hard to see how enterprises could justify doubling their per-seat spending on endpoint security just to cover gaping holes in their existing, signature-based protection model.

Now we hear serious talk of OEM relationships forming with seemingly competitive incumbent anti-malware players looking for lightweight endpoint agents that can be pushed out to protect a wide variety of endpoints in sensitive contexts, such as e-commerce, finance and online banking. Trend Micro's partnership with BigFix is part of this. We expect similar deals from some of Trend's second-tier competitors, as well as organic initiatives to broaden the threat detection and blocking story to encompass software distribution, configuration management and compliance.

Similarly, we see opportunities for the second (or is it third?) wave of DLP-focused M&A as attention shifts from endpoint protection to 'mission assurance,' compliance and business continuity in the context of 'assumed owned' endpoints. Trustwave's recent acquisition of data encryption and IRM vendor BitArmor suggests one model for (managed) endpoint security in the age of APTs: tight integration between endpoint monitoring, DLP and data-level encryption and policy enforcement. For Trustwave, the driving impetus is the compliance requirements of PCI, HIPAA and so on.

But BitArmor's technology – married to DLP from Vericept and access control from NAC appliance vendor Mirage, along with a mature management infrastructure – provides the elements necessary for a data-centric (as opposed to endpoint-centric) security model that could be one alternative to the threat of APTs and 'assumed owned' environments. The approach provides data-level encryption and granular controls over data flows that take into account the context of the data access: the role and identity, asset information, internal and external reputations, and threat intelligence, and so on.

BitArmor's good fortune could translate into interest in competitors like Verdasys and BeCrypt. Verdasys' Enterprise Information Protection architecture seems increasingly relevant to this conversation, comprising data discovery, classification and inventory; data-level access management; information usage management or authorization; and forensics and reporting (for compliance and risk management). Verdasys' Digital Guardian platform has won kudos for its flexibility, but complexity and cost to implement make it a tough sell in the enterprise space. Still, the company's technology could be a good fit for competing MSSPs, SIs or a larger security, systems management or PCLM vendor in search of an endpoint enforcement and leak-protection story.

Looking more broadly at the fast-converging worlds of endpoint security and endpoint management around a holistic PCLM story, we see lots of opportunity for both partnerships and M&A, with compliance and cybercrime and APTs as drivers. *Dell's recent acquisition of KACE* points to the burgeoning market for quick-and-dirty appliance-

based systems management for the SMB and gives Dell a platform on which to build its endpoint security offerings. In the software space, Sophos' pending IPO could give it the liquidity necessary to make a play for an established patch and configuration management vendor. If it wants to jump the fence from consumers and SMBs to large enterprises, Kaspersky Lab also will need to buy or build a management platform that can make it credible.

BigFix finds itself on the wrong end of IBM's recent shake-up of its ISS division, losing one of its larger PC management partners and putting more weight on its partnership with Trend. A take-out or long-rumored IPO by BigFix could, in turn, leave Trend exposed. Vendors like Triumfant have an increasingly relevant story around aggregated threat intelligence married to signature-less endpoint monitoring with elements of configuration management that offer endpoint-specific (rather than generic) remediations to problems. It says Lumension Security has sniffed around, but it has had little interaction with other vendors in that space. Otherwise, look for players like Shavlik Technologies and Lumension to be in play for larger vendors.

5.6 BOTS ARE HOT (AGAIN)

We believe that botnet protection firms also stand to benefit from broad demand for better threat intelligence. We have not been shy about pointing out the (obvious) challenges such firms face in marketing their technology and services to an enterprise marketplace. Today, rank-and-file enterprises still don't consider bots worthy of a separate investment beyond their existing endpoint and gateway anti-malware protections. Concern about spotting bot command-and-control networks and data exfiltration is still limited to the upper end of the enterprise market, and to sensitive verticals like finance, banking, pharmaceuticals, high tech and government.

That has been a challenge for pure-play anti-botnet providers like Damballa and FireEye as they *look to broaden the appeal* of their products beyond the high end of the enterprise market. The good news for these vendors is that their specialty – detecting and tracking botnet command-and-control networks, their supporting infrastructure and compromised hosts – has become the common ingredient of an increasing number of security 'secret sauces.' We see opportunities for take-outs by ISVs looking to roll intelligence about botnet command-and-control networks and reputation monitoring into native security-monitoring offerings, or to bolster vulnerability and threat research.

Botnets, overall, are just one flavor of malware, and incumbent anti-malware vendors tell us they have already developed this capability natively, leveraging large consumer and enterprise installed bases and global threat research organizations. (We note that Symantec's *Mi5 acquisition* brought anti-botnet capabilities in-house.) We don't doubt that incumbents have wrapped their arms around botnet malware, but we're unconvinced they have a firm grasp on the ever-evolving command-and-control infrastructure that the threats use, and think there's competitive differentiation around spotting

fast flux networks or picking out the telltale signs of nascent botnets from analysis of DNS activity, domain registrations and the like. That could translate into targeted acquisitions by diversified anti-malware vendors or MSSPs that want to boost their threat analytics capabilities.

Similarly, secure gateway and Web threat protection vendors like Blue Coat and Websense could get an immediate boost from the kinds of traffic analysis, threat intelligence and reputation information that firms Damballa and FireEye hold, especially given the intersection of bots and Web-based attacks and communications. For the right price, second-tier anti-malware firms like Kaspersky Lab or Sophos could scoop up an anti-botnet specialist.

We also see opportunities for focused MSSPs like Perimeter eSecurity, Trustwave and the like to upsell anti-botnet services to their customer bases – succeeding where pure-play vendors have often fallen short. Finally, firms like NetWitness, Solera and Niksun could partner or acquire botnet-detection IP to supplement their traffic monitoring and detection capabilities.

SECTION 6

Vendor Profiles

6.1 ARCSIGHT

Company at a glance

ArcSight sells enterprise security information management systems as well as a line of low-end log management appliances. Recently, it has begun integrating unstructured application and user data to those platforms, generating visibility into patterns of activity over time.

LOCATION	5 Results Way, Cupertino, California 95014
FOUNDED	2000
KEY EXECUTIVES	Tom Reilly, CEO; Hugh Njemanze, CTO; Stewart Grierson, CFO
CUSTOMERS	725 (FY09)
SALES RATIO	overall = 60% direct, 40% indirect (Feb.09)
REVENUE	\$136m (FY09)
US GAAP PROFITABLE	Yes (FY09)
CASH-FLOW POSITIVE	Yes (FY09)
EMPLOYEES	400 (FY09)
TOTAL FUNDING RAISED	\$28.5m
ROUNDS/INVESTORS	Series C: 11/2002, \$3m from In-Q-Tel; Series B: 4/2002, \$9.5m from Kleiner Perkins Caufield & Byers, Silicon Valley Internet Capital, Integral Capital Partners; Series A: 5/2000, \$16m from Silicon Valley Internet Capital

The 451 Take

ArcSight is perhaps the most visible player at the high end of the spectrum for log management and SIEM, where it competes against EMC/RSA, IBM Tivoli and CA as well as Attachmate/NetIQ. A range of smaller players compete as well, including Inspekt Security, Intellitactics, LogRhythm, netForensics, NitroSecurity, Novell, OpenService, Q1 Labs, SenSage, Symantec, Tier-3, eIQNetworks, TriGeo and Cisco's CS-MARS. The past two years have been good to ArcSight, with demand for its products driven both by compliance and a more complex threat landscape. More recently, the company has made steps to move downmarket. It has also sought to broaden the scope of monitoring, adding unstructured data to the Logger log management product and focusing on more complex threats with the FraudView product. There's interest in what ArcSight has to offer from a wide range of larger players, and we've heard rumors of talks to acquire the company; 2010 should be an interesting year for ArcSight.

6.2 BRIGHTCLOUD

Company at a glance

BrightCloud offers a hosted reputation management service via its URL database and filtering service.

LOCATION	4810 Eastgate Mall, San Diego, California 92121
FOUNDED	2005
KEY EXECUTIVES	Quinn Curtis, president & CEO; Chris Harris, categorization architect; Ron Hegli, member, technical advisory board; Hal Lonas, VP of engineering
SALES RATIO	100% direct
US GAAP PROFITABLE	Yes
CASH-FLOW POSITIVE	Yes
AVERAGE DEAL SIZE	more than \$150,000
DEAL RANGE	\$50,000 to more than \$1m
EMPLOYEES	10 FTE, 60 contracted: 80% in tech
TOTAL FUNDING RAISED	\$2.6m
ROUNDS/INVESTORS	Series B: 2008, \$1.2m (in venture debt funding) from Apex Venture Partners; Series A: 2006, \$1.4m from founders

The 451 Take

With Web-based threats the new battleground for enterprise endpoints, BrightCloud finds itself with a service that a lot of other companies need: broad-spectrum Web reputation data and analytics. The key here is BrightCloud's robust Web crawler and AI that allow the company to identify and categorize malicious or compromised websites. For now, BrightCloud is content to be an 'arms dealer,' licensing its services to a wide range of OEM partners in the endpoint and network protection business. The company is also broadening its reach from Web to other kinds of traffic, such as IM, botnets and so on. But BrightCloud could easily be subsumed as larger security firms look for more and better data on malicious URLs, IP addresses and the like to inform their policy engines.

6.3 COMMTOUCH SOFTWARE

Company at a glance

Commtouch offers a variety of Web and messaging security products; including anti-spam, anti-botnet, URL-filtering and email-reputation services.

LOCATION	4A Hatzoran St, PO Box 8511, Netanya 42504, Israel
FOUNDED	1991
KEY EXECUTIVES	Gideon Mantel, CEO; Amir Lev, CTO; Ido Hadari, COO; Ron Ela, CFO
REVENUE	\$14m (FY08)
US GAAP PROFITABLE	Yes (FY08)
CASH-FLOW POSITIVE	Yes (FY08)
EMPLOYEES	69; 53 located in Israel (FY08)

The 451 Take

Commtouch has built a thriving products and services business by leveraging relationships with top-tier ISPs, becoming a sought-after arms dealer to security hardware and software vendors with its anti-spam, anti-botnet and Web security services. The company has been adding OEM partners at a breakneck pace. But with services still focused on messaging and Web and plenty of competitors in both areas, we wonder what Commtouch's next move will be. The company has mentioned email archiving and encryption as possible areas of expansion. We'll see whether it makes good on its promises in 2010, or goes in a different direction – perhaps adding brand protection and reputation monitoring to the mix.

6.4 HBGARY

Company at a glance

HBGary sells several tools for malware and forensic analysis, as well as incident response.

LOCATION	1029 H Street, Suite 308, Sacramento, California 95814
FOUNDED	2003
KEY EXECUTIVES	Greg Hoglund, CEO; Rich Cummings, CTO
REFERENCE CUSTOMERS	Customers include consulting firms, manufacturing (high tech or any vertical where IP theft is an issue), banking, and pharmaceutical companies
CUSTOMERS	250
SALES RATIO	80% direct
US GAAP PROFITABLE	Yes
CASH-FLOW POSITIVE	Yes
AVERAGE DEAL SIZE	varies by product line: Responder PRO = \$35,000; Digital DNA = \$25,000
EMPLOYEES	21 (17 in technical positions)
TOTAL FUNDING RAISED	No outside funding

The 451 Take

HBGary is a six-year-old firm that is the brainchild of CEO, founder and rootkit expert Greg Hoglund. HBG arose in 2003, building on research Hoglund had done in rootkit detection and forensics. The core offering here is malware forensics that are qualitatively different from what's available from commercial anti-malware research labs. That doesn't sound like much, but HBGary's roster of federal customers, integration with McAfee's ePO and partnerships with up-and-comers in the cyberintelligence community (like Palantir) suggest that people are interested in what HBG has to offer.

6.5 MARKMONITOR

Company at a glance

MarkMonitor offers anti-fraud technology (including anti-malware and anti-phishing) as part of an overall brand management strategy.

LOCATION	303 Second Street, Suite 800N, San Francisco, California 94107
FOUNDED	1997
KEY EXECUTIVES	Irfan Salim, president & CEO; Paul Dagum, CTO; Ihab Shraim, CSO & VP of network and system engineering; Tom Ryden, SVP, finance
REFERENCE CUSTOMERS	American Apparel, Comerica Bank, Du Pont, FedEx, Sovereign Bank, Time Warner, Under Armour, Liberty Mutual, World Wrestling Entertainment, UBS
CUSTOMERS	800
SALES RATIO	85% direct
DEAL RANGE	\$25,000 - \$1,500,000
EMPLOYEES	220; 160 in tech
TOTAL FUNDING RAISED	\$68m (company claims its last-round funds are still 'in the bank')
ROUNDS/INVESTORS	Series E: 06/2006, \$24m from Polaris Venture Partners, Institutional Venture Partners, Cargill Ventures, Focus Ventures, Foundation Capital; Series D: 04/2006, \$12m from Cargill, Focus, Foundation, Institutional VP; Series C: 12/2003, \$20m from Focus, Foundation, Institutional VP; Series B: 12/2002, \$8m from Foundation; Series A: 11/2002, \$4m from founders and angel investors

The 451 Take

Like Cyveillance, MarkMonitor has found a sweet spot combining anti-phishing services with more specialized reputation- and brand-monitoring offerings. We think there are broad applications for that kind of capability as part of a larger threat protection and compliance story. MarkMonitor claims it still has cash on hand from its last round of funding in 2006. We'll be interested to see how the company uses that to expand its offerings in 2010, or whether larger firms come looking for a piece of what the company has to offer.

6.6 MEMENTO SECURITY

Company at a glance

Memento Security provides fraud prevention and detection in multiple areas, including check fraud, employee fraud and healthcare fraud.

LOCATION	55 Network Drive, Burlington, Massachusetts 01803
FOUNDED	2003
KEY EXECUTIVES	BC Krishna, founder & CEO; Mike Braatz, VP of marketing; Joe Walsh, chief software architect; Paul Whitelam, VP of product management
SALES RATIO	overall: 90% direct, 10% indirect (Dec 09)
EMPLOYEES	70 total; 60% in tech positions (Dec 09)
TOTAL FUNDING RAISED	\$22m
ROUNDS/INVESTORS	Series D: 11/2008, \$10m from Bain Capital Ventures, .406 Ventures, Rock Maple Ventures; Series C: 4/2007, \$7.5m from .406, Bain, Rock Maple, private investors; Series B: 4/2005, \$3m from Bain, Rock Maple, private investors; Series A: 2004, \$1.5m from Rock Maple, private investors

The 451 Take

Tiny Memento Security has carved out a niche in the much-larger market for fraud monitoring. The company's artificial intelligence is particularly strong at stopping new forms of complex banking and financial fraud schemes such as 'third-party credit card bust-outs,' leveraging source data from a broad spectrum of structured and unstructured data from core processing applications, and picking out patterns of unusual activity between groups of individuals. The focus now is on banking fraud, but we (and Memento) see broader applications for its analytics and anti-fraud IP in the enterprise context.

6.7 NETWITNESS

Company at a glance

NetWitness sells network forensics technology and risk management software to customers in verticals such as government, defense, law enforcement, banking and critical infrastructure.

LOCATION	500 Grove Street, Suite 300, Herndon, Virginia 20170
FOUNDED	2006
KEY EXECUTIVES	Amit Yoran, CEO; Nick Lantuh, president; Tim Belcher, CTO; Edward Schwartz, CSO; Dana Duffy, CFO
CUSTOMERS	100
REVENUE	N/A
US GAAP PROFITABLE	N/A
CASH-FLOW POSITIVE	N/A
EMPLOYEES	72
TOTAL FUNDING RAISED	\$7.5m
ROUNDS/INVESTORS	founder Amit Yoran, Summit Partners

The 451 Take

NetWitness CEO and founder Amit Yoran has a penchant for spotting trends in the security industry. This was the case when he bought the underlying technology for NetWitness' network forensics platform from defense contractor ManTech. Today, the company's products occupy an increasingly important niche: aggregating, correlating and reconstructing disparate data from network infrastructure and endpoints in a way that allows security investigators to spot sophisticated attacks, insider threats, network performance and compliance-related issues. This stuff isn't for the faint of heart, and NetWitness' audience so far is large enterprise and government agencies with the talent and resources to use its platform. But we see a much broader market for what NetWitness does, especially as it builds bridges with a broader ecosystem of MSSPs, ISVs and threat intelligence vendors.

6.8 NICE

Company at a glance

NICE Systems provides data analytics products and services that monitor, record, analyze and log transactions and interactions for banking, financial services, retail and other verticals.

LOCATION	8 Hapnina St, PO Box 690, Ra'anana 43107, Israel
FOUNDED	1986
KEY EXECUTIVES	Zeev Bregman, CEO; Dafna Gruber, CFO; Israel Livnat, corporate VP and president, security group
CUSTOMERS	more than 24,000
REVENUE	\$624m (FY08)
US GAAP PROFITABLE	Yes (FY08)

The 451 Take

NICE has been steadily moving upmarket in recent years, as it used acquisitions to broaden its offerings from call-center monitoring to address other types of fraud. Most recently, the company picked up anti-money-laundering firm Fortent in a deal that sets the stage for closer integration between the now-separate worlds of consumer-based anti-fraud technology and higher-end services that can detect intrabank activity and sophisticated transactions. There's a bigger GRC story here, as NICE pulls more data into its monitoring platform, including VoIP and mobile telephony, etc.

6.9 PALANTIR TECHNOLOGIES

Company at a glance

Founded by a handful of PayPal alums, Palantir offers information analytics to government and financial verticals.

LOCATION	100 Hamilton Ave, Palo Alto, California 94301
FOUNDED	2004
KEY EXECUTIVES	Alex Karp: CEO; Nathan Gettings, CTO
SALES RATIO	100% direct
AVERAGE DEAL SIZE	\$1,500,000
EMPLOYEES	215; tech to non-tech is 3:1
TOTAL FUNDING RAISED	Undisclosed
ROUNDS/INVESTORS	Venture capital comes from The Founders Fund, In-Q-Tel and Reed Elsevier Ventures. How much was raised and in how many rounds remains undisclosed.

The 451 Take

As enterprises and governments gear up to battle APTs, platforms like Palantir's analytics become increasingly important, allowing security analysts to integrate disparate structured and unstructured data – security events, threat forensics, message traffic, netflows, user access and geospatial data. The object is to map the relationships between the discrete events and pieces of data that together make up sophisticated attacks like GhostNet and Aurora. For now, Palantir's products are sold to a small audience: government and financial services that are most exposed to the APT problem. That audience could broaden, especially for a managed service, as more organizations try to get their arms around stealthy, multivector attacks, network compromises and the theft of sensitive data.

6.10 QINETIQ

Company at a glance

QinetiQ is a London-based research and defense contracting firm with origins in the British government. In 2009, the company bought Cyveillance, a vendor that gathers intelligence about the illegal use of digital content, and the misrepresentation and resale of digital content into real-world assets.

LOCATION	85 Buckingham Gate, London SW1E 6PD, UK
FOUNDED	2001
KEY EXECUTIVES	Leo Quinn, CEO; David Mellows, CFO
REVENUE	£1.6bn (\$2.3bn in FY09)
US GAAP PROFITABLE	Yes (FY09)
CASH-FLOW POSITIVE	Yes (FY09)
EMPLOYEES	more than 14,060 (FY09)
TOTAL FUNDING RAISED	£42m (\$68m in 2/03)
ROUNDS/INVESTORS	Series A: 2/2003, £42m (\$68m) from The Carlyle Group

The 451 Take

Government-focused MSSP QinetiQ made a bet on cybercrime intelligence when it picked up Cyveillance, but now it needs to find a way to marry its acquired Web-crawling, anti-phishing, malware-analysis and brand-monitoring chops with QinetiQ's mostly traditional managed security offerings and services. QinetiQ has a deep pool of talent and technology for hardening enterprise networks and providing quality threat and malware intelligence. We think that could be a potent mixture for both government and enterprise, but there are lots of blank areas that need filling in on QinetiQ's roadmap.

6.11 RSA (THE SECURITY DIVISION OF EMC)

Company at a glance

RSA is the security division of storage giant EMC. It makes software and hardware for IAM, encryption, log management, enterprise security information management, anti-fraud, and ADL and data-classification.

LOCATION	174 Middlesex Turnpike, Bedford, Massachusetts 01730
FOUNDED	1986
KEY EXECUTIVES	Arthur Coviello, EVP of EMC & president of RSA; Bret Hartman, CTO; Christopher Young, SVP, products
REFERENCE CUSTOMERS	AT&T, Credit Suisse, Lufthansa, Oracle, General Electric, Rolls-Royce, Giant Eagle, Hershey's, Visa International, Nintendo Co
REVENUE	\$581m (RSA, FY08); \$14.8bn (EMC, FY08)
US GAAP PROFITABLE	Yes (EMC, FY08)
CASH-FLOW POSITIVE	Yes (EMC, FY08)
EMPLOYEES	42,100 (EMC, FY08)

The 451 Take

EMC's security division, RSA, has begun to integrate the key elements of its platform: DLP, risk-based authentication and the Ionix IT management platform, as well as EnVision SIEM. The recent addition of the GRC platform it acquired with Archer Technologies will be key, as EMC/RSA plans on reaching a broader enterprise audience with a process management tier that integrates information from management silos and enriching it with metadata. The outcome is more meaningful data that can be tied to remediation and resolution actions.

6.12 TEAM CYMRU

Company at a glance

Team Cymru is a not-for-profit consulting group that sells threat intelligence feeds around malware, phishing and botnet command-and-control networks.

LOCATION	16W361 S. Frontage Road, Suite 100, Burr Ridge, Illinois 60527
FOUNDED	1998, incorporated in 2004
KEY EXECUTIVES	Rob Thomas, CEO; Jeff Vosburg, COO; Dave Deitrich: CTO; Jerry Martin, CFO; Dave Munsun, CIO

The 451 Take

Cymru, which employs 32 people, still won't say much about its activities. What we do know puts Cymru firmly within the ranks of boutique threat-intelligence firms. The company leverages relationships with ISP partners globally, harvesting terabytes of information on malicious traffic and attacks. It combines that with up-close surveillance of cybercriminal networks. The threat intelligence it derives from that is shared with law enforcement and sold to Cymru's customers. As attacks become vertical-specific and target-specific, outfits like Cymru become increasingly relevant: offering differentiated threat intelligence on emerging attacks, brand monitoring and the like. The question is whether Cymru and others can find a way to scale their operations and grow out of their niche.

6.13 TRIUMFANT

Company at a glance

Triumfant's products offer endpoint-health monitoring that enables companies to detect attacks on or changes to systems on their network without signatures, then remediate unauthorized changes on the fly.

LOCATION	800 King Farm Boulevard, Rockville, Maryland 20850
FOUNDED	2002
KEY EXECUTIVES	John Prisco, president & CEO; David Hooks, founder & CTO; Jim Ivers, CMO
CUSTOMERS	20; 100,000 seats
SALES RATIO	80% direct
US GAAP PROFITABLE	No
CASH-FLOW POSITIVE	No
AVERAGE DEAL SIZE:	\$120,000
EMPLOYEES	less than 20; 50% in tech
TOTAL FUNDING RAISED	\$19.5m
ROUNDS/INVESTORS	Series C1: 4/2006, \$3m from Novak Biddle Venture Partners; Series C: 3/2006, \$6m from Core Capital Partners, Anthem Capital Management, Inflection Point Ventures; Series B: 3/2005, \$9m from Core, Anthem, Inflection Point; Series A: 6/2004, \$1.5m from MCNC Ventures, Tri-State Investment Group IV

The 451 Take

We're intrigued by Triumfant's approach to threat prevention, which combines change and configuration management with a novel take on application monitoring and whitelisting. The company addresses the pain points that government agencies and IP-rich companies are suffering from the need to respond to zero-day threats and understand the compliance impact of infections. Triumfant's story is upbeat, but we note a flat customer count from Q2 2009. Growing its enterprise business will be key in 2010. If Triumfant's value is really around attack impact analysis and change management, we're interested in what other partners it can scare up. Anti-malware, SIEM and GRC are all options.

6.14 TRUSTEER

Company at a glance

Trusteer's Rapport product secures data sent back and forth in high-value Web sessions and transactions by thwarting malicious programs such as rootkits that are plaguing customers in its key vertical of online banking and brokerage.

LOCATION	13 Noach Mozes St, Tel Aviv 67442, Israel
FOUNDED	2006
KEY EXECUTIVES	Mickey Boodaei, CEO; Amit Klein, CTO; Shmulik Regev, VP & chief architect
REFERENCE CUSTOMERS	ING, The Royal Bank of Scotland, Alliance & Leicester, SVB Financial Group, BBVA Compass Bank
CUSTOMERS	more than 50
SALES RATIO	100% direct
US GAAP PROFITABLE	No
CASH-FLOW POSITIVE	Yes
DEAL RANGE	\$50,000 – “several million” per year
EMPLOYEES	35; 30 in tech positions
TOTAL FUNDING RAISED	\$10m
ROUNDS/INVESTORS	Series B: 10/2008, \$6m from U.S. Venture Partners; Series A: 7/2006, \$4m from Shlomo Kramer (of Imperva, Check Point) and other private investors

The 451 Take

Like close competitor Verdasys, Trusteer is positioning its technology to play in a world where pervasive malware has made it impossible to ensure the integrity of endpoints. Rather than trying to keep every last bit of malicious code off endpoints, Trusteer focuses on protecting data itself by taking a position low down in the OS kernel and locking down sensitive transactions such as online banking sessions, from keystroke through delivery to the Web application. Given the poor state of protection offered by legacy products and the growing frustration with them, we think Trusteer's Rapport offering will be very attractive to C-level executives in regulated industries like banking, financial services, healthcare and government. But as we've said before, writing a rootkit that plays nicely with Windows and other programs is no easy task. Given that Rapport is just a rootkit wearing a white hat, can Trusteer guarantee that it won't clash with other software on the endpoint, blue-screen Windows desktops, or find itself locked out by malware running even lower on the host than it is? Time will tell.

INDEX OF COMPANIES

3FN	33	Blue Ridge Networks	42
.406 Ventures	66	BrandProtect	54
Adobe	1, 11, 13, 14, 16, 37, 38	BrightCloud	62
Agilance	57	BT	55
Akamai	8	CA	46, 52, 57, 61
Alert Logic	51, 57	Card Systems Inc	23
Altiris	38, 39	Cargill Ventures	65
Anthem Capital Management	73	Cassandra Security	15, 45, 55
Apex Venture Partners	62	Check Point	74
Arbor	8	ChoicePoint	1, 23
Archer Technologies	57, 71	Cisco Systems	43, 55, 56, 61
ArcSight	48, 50, 51, 56, 57, 61	Citadel	38
ATT	44	Citibank	22
Attachmate/NetIQ	61	CommTouch Software	63
Authentium	42, 57	Core Capital Partners	73
Bain Capital Ventures	66	CoreTrace	40
Barracuda	55	Cyveillance	19, 43, 44, 45, 54, 65, 70
BeCrypt	42, 58	Damballa	43, 44, 59, 60
BigFix	38, 39, 58, 59	Dave & Busters	1
Bit9	40, 49	eIQnetworks	51, 56
BitArmor	53, 58	EMC/RSA	44, 49, 50, 51, 57, 61, 71
Blue Coat Systems	44, 60	Envisional	54

Eurekify 46	Inflection Point Ventures 73
F5 55	In-Q-Tel 32, 61, 69
Fidelis Security Systems 48, 53	Inspekt Security 61
FireEye 43, 44, 59, 60	Institutional Venture Partners 65
Fluke Networks 47	Integral Capital Partners 61
Focus Ventures 65	Intel 1, 37
Forrester 24, 25	Intellitactics 51, 56, 57, 61
Fortent 68	IPLocks 24
Foundation Capital 65	Ironport 43
Gartner 24	ISS 47, 59
GigaTrust 53	Juniper Networks 13, 46
Google 1, 13, 14, 17, 28, 37, 38, 49	Kaspersky 59, 60
Guardian Analytics 42, 46	Kleiner Perkins Caufield & Byers 61
Hannaford Supermarkets 1, 2, 23	Liquid Machines 53
HBGary 49, 55, 64	LogLogic 51, 56
Heartland Payment Systems 1, 2, 6, 23, 31	LogRhythm 56, 61
IBM 42, 46, 47, 59, 61	Lookingglass Cyber Solutions 51
iDefense 14, 49	Lumension Security 38, 39, 59
Imperva 74	Mandiant 49, 55
InDorse Technologies 53	ManTech 67
Industrial Defender 19, 20	MarkMonitor 43, 45, 54, 65

McAfee 9, 19, 23, 28, 33, 38, 39, 40, 42, 43, 47, 49, 51, 55, 56, 64	Orchestria 52
McColo 33	Palantir 51, 52, 55, 64, 69
MCNC Ventures 73	Palo Alto Networks 48
Memento Security 46, 47, 66	Perimeter eSecurity 57, 60
Microsoft 5, 9, 10, 14, 16, 23, 25, 40, 42, 49	Persysent Software 39
Mirage 58	Polaris Venture Partners 65
Modulo Security 57	Ponemon Institute 24, 25
Morgan Stanley 1	Preventsys 38
netForensics 51, 61	Prevx 41, 57
NetScout 47	Prolexic 44
NetWitness 27, 28, 47, 48, 51, 60, 67	Q1 Labs 57, 61
NICE Systems 68	QinetiQ 19, 43, 44, 54, 70
Niksun 47, 48, 60	Quest 57
NitroSecurity 51, 56, 61	Raytheon 19
Norman ASA 38	RBS 22
Northrop Grumman 13, 14	RedCannon Security 42
Novak Biddle Venture Partners 73	Safend 53
Novell 61	Savant Protection 40
Oakley 19	Secure Computing 43
OfficeMax 1, 23	SecureWorks 44, 49, 55
OpenService 61	SenSage 56, 61
	Shavlik 38, 59

Shavlik Technologies 38, 59	TJX 1, 6, 22, 23, 24, 25
Signacert 40	Trend Micro 9, 11, 38, 39, 42, 58
Silicon Valley Internet Capital 61	TriGeo 51, 56, 61
Skybox 57	Tripwire 39
Skybox Security 57	Tri-State Investment Group 73
Solera 48, 51, 60	Triumfant 39, 59, 73
Solera Networks 48	Trusteer 41, 42, 57, 74
Solidcore 19, 40	Trustwave 44, 53, 55, 57, 58, 60, 70
Sophos 11, 60	U.S. Venture Partners 74
Sourcefire 48	Varonis 53
Splunk 51, 56	Verdasys 53, 58, 74
Summit Partners 47, 67	Vericept 53, 58
Sunbelt Software 38	VeriSign 14, 44, 49
Symantec 9, 11, 13, 27, 28, 33, 38, 39, 42, 48, 56, 59, 61	Verizon Business 25, 55
Team Cymru 45, 72	Vigilant 57
Telus 57	Vontu 48
Tenable 51, 56	WildPackets 47
The Carlyle Group 70	Wireshark 47
The Founders Fund 69	ZafeSoft 53
The Wall Street Journal 27	
Tier-3 61	