InformationWeek
:: reports

Next

Reports.InformationWeek.com February 2014 $99

Advanced Attacks Demand New Defenses

Cyber-criminals wielding APTs have plenty of innovative techniques to evade network and endpoint defenses. It's scary stuff, and ignorance is definitely not bliss. How to fight back? Think security that's distributed, stratified, and adaptive.

By Kurt Marko

Report ID: S7780214

TABLE OF CONTENTS

ABOUT US

InformationWeek Reports' analysts arm business technology decision-makers with real-world perspective based on qualitative and quantitative research, business and technology assessment and planning tools, and adoption best practices gleaned from experience.

OUR STAFF

Lorna Garey, content director; lorna.garey@ubm.com
Heather Vallis, managing editor, research; heather.vallis@ubm.com
Elizabeth Chodak, copy chief; elizabeth.chodak@ubm.com
Tara DeFilippo, associate art director; tara.defilippo@ubm.com

Find all of our reports at reports.informationweek.com.

Kurt Marko
InformationWeek Reports

Kurt Marko is an InformationWeek and Network Computing contributor and IT industry veteran, pursuing his passion for communications after a varied career that has spanned virtually the entire high-tech food chain, from chips to systems. Upon graduating from Stanford University with a BS and MS in electrical engineering, Kurt spent several years as a semiconductor device physicist, doing process design, modeling, and testing. He then joined AT&T Bell Laboratories as a memory chip designer and CAD and simulation developer.

Moving to Hewlett-Packard, Kurt started in the laser printer R&D lab doing electrophotography development, for which he earned a patent, but his love of computers eventually led him to join HP's nascent technical IT group. He spent 15 years as an IT engineer and was a lead architect for several enterprise-wide infrastructure projects at HP, including the Windows domain infrastructure, remote access service, Exchange email infrastructure, and managed web services.

EXECUTIVE SUMMARY

Security threat and response is a vicious circle of escalating (and increasingly cagey) attacks and sophisticated (and increasingly costly) defenses. The latest generation of malware includes deviously creative evasive techniques crafted to exploit ambiguities in the Internet's underlying technology, flaws in network software stacks, and limitations of security appliances.

One example operates at the network-protocol level to bypass firewalls and intrusion-prevention systems by hiding malicious traffic within abnormal, but still compliant, TCP/IP packets. Another category works entirely within common applications using normal rules for web traffic. They don't so much trick network security software as bypass it using HTML5 and embedded scripts to distribute malicious payloads. In this report, we discuss these techniques, how IT teams can test their level of exposure, and how to detect and block attacks using advanced packet normalization.

The Goal: Penetration. The Tactic: Trickery

Cyber-security is often characterized as a game of cat and mouse. If so, recent history reads like a Tom and Jerry cartoon — the mice seem to be winning. We all know the problem: Our adversaries are organized cybercrime enterprises or nation-states with plenty of time and money. There's also the inherent truth of every security mission — the enemy has to be right only once. And today, the enemy seems to control the rules to the game.

Among the most dangerous new gambits are attacks using so-called advanced evasion techniques. It's a handy way to classify attacks designed to trick and bypass layered defenses incorporating perimeter firewalls, IDS/IPS implementations, web content filters, and endpoint protection software. Each layer is designed to block delivery of an exploit that targets underlying vulnerabilities, whether open server ports, unpatched OS holes, or trusting end users.

Attackers are looking to hide malicious code from network security layers like Jerry slips by Tom's claws, and infosec teams need to pay attention. We agree with the three primary security challenges for the coming year that Cisco laid out in its 2014 Annual Security Report:

>> **Greater attack surface area:** The vanishing network border with the addition of cloud services and external mobile users is colliding with the increasing value of closely held digital assets, whether corporate intellectual property or customer personal and financial information.

>> **Proliferation and sophistication of the attack model:** Advanced persistent threats can lurk inside dozens or hundreds of compromised systems.

>> **Complexity of threats and solutions:** The attacker's ultimate goal is almost always

Figure 1

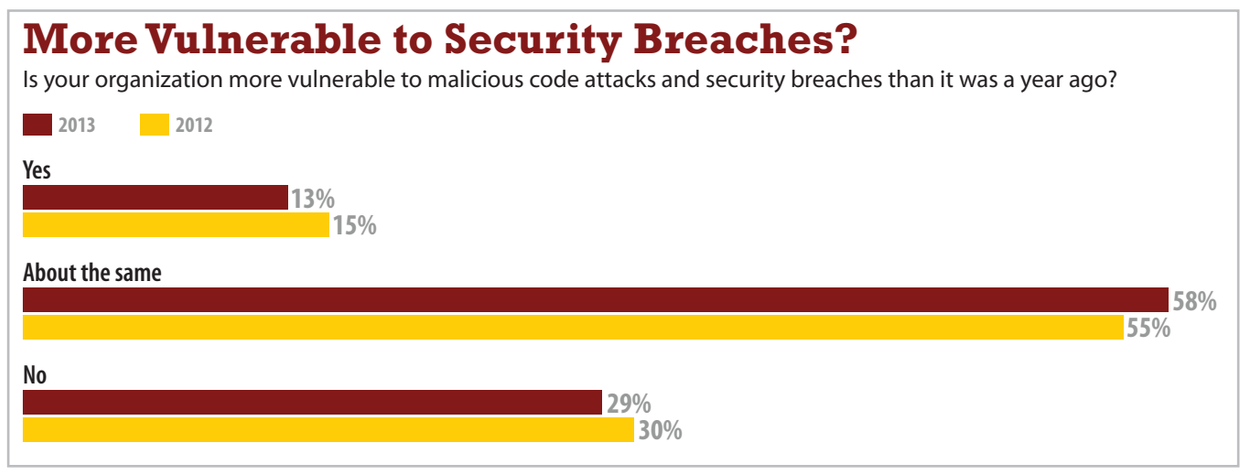

Base: 1,029 respondents in March 2013 and 946 in March 2012 R6820513/2
Data: InformationWeek Strategic Security Survey of business technology and security professionals at organizations with 100 or more employees

penetration of the datacenter, where bulk collection of sensitive and commercially valuable information is far more efficient than attacking individual client devices. But this usually entails exploiting many systems, breaching several network perimeters, and planting an assortment of persistent malware (data collectors, intermediate repositories, command-and-control systems) on compromised systems.

The key: exploitation of trusted systems and applications.

"Attackers are coming up with new methods for embedding their malware in networks, remaining undetected for long periods, and stealing data or disrupting critical systems," write the Cisco report authors. "Using methods ranging from the socially engineered theft of passwords and credentials to stealthy, hide-in-plain-sight infiltrations that execute in minutes, malicious actors continue to exploit public trust to effect harmful consequences."

To get into datacenters, such threats may exploit ambiguities in network protocol standards, concomitant variances in how network stacks are implemented, and limitations in the packet inspection capabilities of most security hardware.

The underlying idea behind advanced evasion techniques — abusing the foundations of network communications — isn't new. As long ago as the late 1980s, researchers

Figure 2

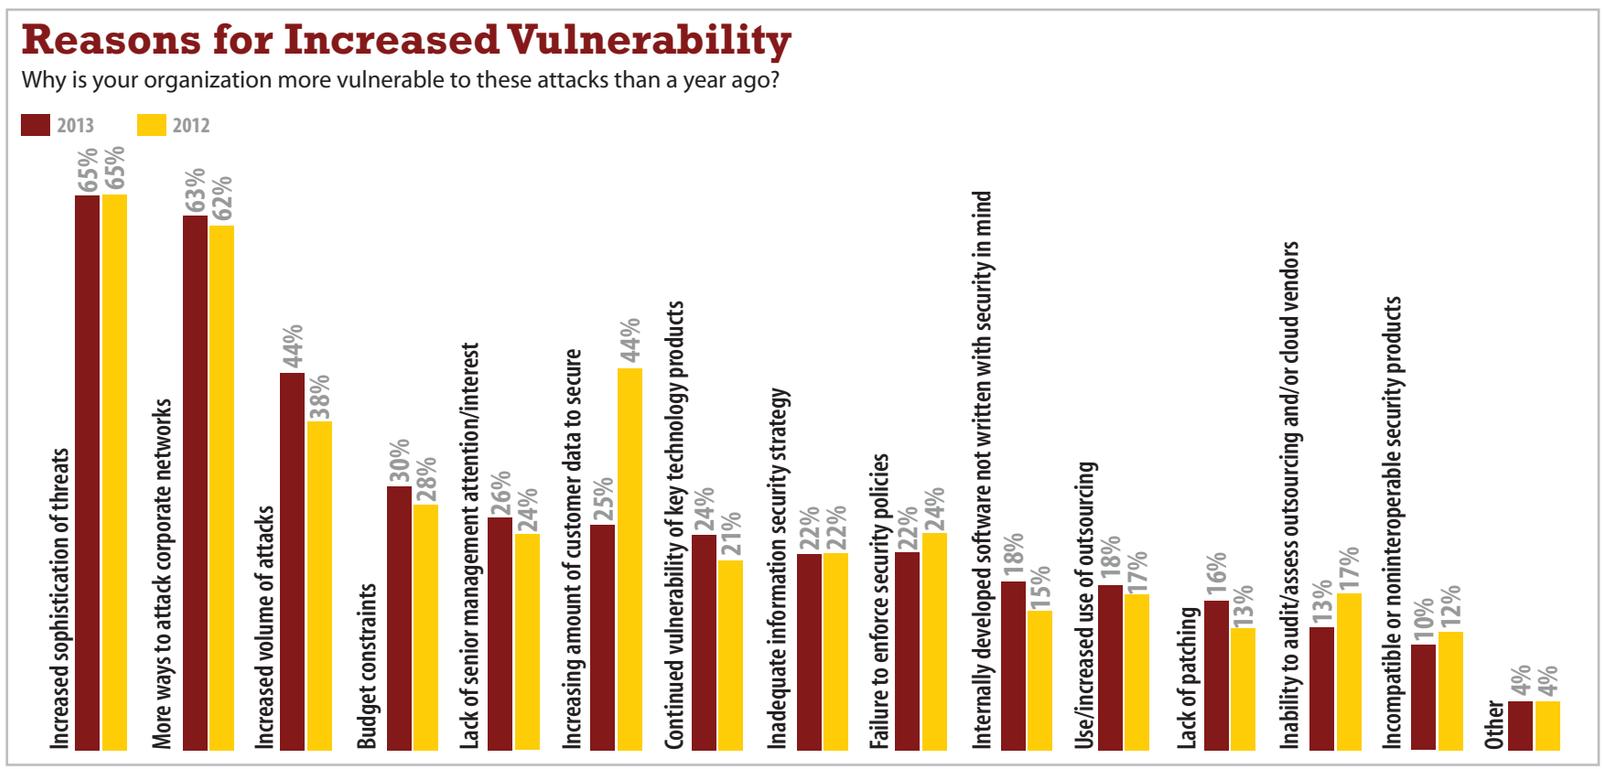

Note: Multiple responses allowed
Base: 135 respondents in March 2013 and 146 in March 2012 at organizations more vulnerable to attacks than a year ago
Data: InformationWeek Strategic Security Survey of business technology and security professionals at organizations with 100 or more employees

explored the notion of burying malicious payloads within seemingly normal TCP/IP connections. The tactic has evolved from academic curiosity to realistic prospect over the past three or four years. While advanced evasion techniques can deliver any variety of malware, because of their complexity, they're commonly used for only the most elaborate and lucrative attacks: APTs.

APTs: Misunderstood At Our Peril

In a recent ISACA survey of 1,500 security professionals, 67.6% of respondents said they're at least familiar with advanced persistent threats. However, 53.4% of those respondents said APTs are similar to traditional threats. They're wrong, and here's why.

The National Institute of Standards and Technology says APTs are characterized by sophisticated levels of expertise and significant resources that allow attackers to achieve their objectives using multiple attack vectors — cyber, physical, and deception. And these objectives aren't just defacing a site or swiping a few credit cards and running. Advanced attackers seek to burrow in so they can exfiltrate information or undermine or impede critical aspects of a mission, program, or organization, not just in the near term but over time.

NIST defines an APT as meeting three criteria: It pursues its objectives repeatedly over an extended period of time. It adapts to defenders' efforts to resist it. And it is determined to maintain the level of interaction needed to execute its objectives.

Item No. 2 is where advanced evasion techniques come in; the two are complementary. APTs are the motivation (extraction of sensitive information) and the attack strategy. Advanced evasion techniques provide access and delivery.

But the NIST APT definition also applies to several other important aspects of advanced evasion methods. They aren't just stealthy, but adaptive and cooperative, often using several techniques working in concert to deliver a malicious payload. Advanced evasion techniques are often used on multiple levels of the network stack simultaneously, meaning one piece of an exploit might be delivered using packet fragmentation and another using flaws in the SMB file-sharing protocol. Similarly, these techniques are adaptive, with the scope and targets of the attack modified over time.

Stonesoft, acquired last year by McAfee, coined the term "AET" and summarizes it in a whitepaper. The gist: Advanced evasion techniques can exist in every protocol and may be combined to create new tactics. The order of combined evasions is important, and the number of possible combinations is huge.

Despite their apparent potential to evade existing network security measures, we could find no published data on how widely advanced evasion techniques are used, the overall attack volumes, or whether use is increasing. To get a sense, most security firms collect and aggregate statistics, logs, and attack signatures from the many devices customers have deployed on their networks. Cisco's SenderBase is one of the largest such repositories, and its reported attack numbers are staggering. Cisco claims it blocks 4.5 billion email and 80 million web requests per

day. It also nixes more than 3,000 endpoint network detections and 50,000 network intrusions per day, with the total number of threat alerts increasing 14% last year.

The bad news for security pros is that most of these are new exploits, not updates of previous threats. McAfee Labs, which operates its own threat monitoring center, said in a recent report that new malware types exceeded 20 million in the third quarter of 2013. Rootkits doubled and signed malware, which poses as legitimate software, rose by almost 50%.

Since advanced evasion techniques are designed to elude such defenses, it's almost impossible to know how many are in the wild. Former Defense Secretary Donald Rumsfeld's famous quip is apropos: "There are known knowns; there are things we know that we know. There are known unknowns; that is to say, there are things that we now know we don't know. But there are also unknown unknowns — there are things we do not know we don't know."

Advanced evasion techniques that security researchers have publicly demonstrated fall into the second category of "known unknowns." However, those in the wild, which could include new and creative implementations, are clearly "unknown unknowns." Getting a good estimate of usage is probably impossible. Still, that doesn't mean we should assume they don't exist.

Ignorance Is Not Bliss

Surprisingly, enterprise IT doesn't always perceive the world as getting more dangerous even as one security researcher says all Fortune 500 companies have been breached. Only 13% of 1,029 security pros responding to the InformationWeek 2013 Strategic Security Survey saw themselves as more vulnerable to malicious code attacks and other security breaches than they were the previous year. However, among that 13%, almost two-thirds cite the increased sophistication of threats and a larger network attack surface as the primary reasons they feel more exposed. We would say their assessment is spot on.

Just 21% of our security survey respondents said they had experienced a breach in the

Figure 3

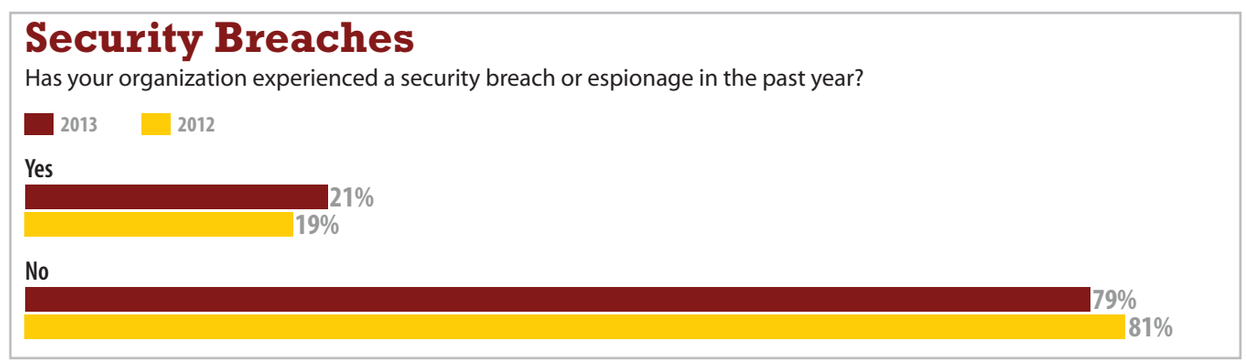

Base: 1,029 respondents in March 2013 and 946 in March 2012 R6820513/4
Data: InformationWeek Strategic Security Survey of business technology and security professionals at organizations with 100 or more employees

past year, up two points since 2012. However, this almost surely understates attack frequency since the most successful exploits are never detected, plus there is always some level of reporting bias on potentially embarrassing questions. Witness the Neiman Marcus incident, which went unnoticed for more than six months, or the Target point-of-sale exploit that was detected only after stolen card numbers started being used en masse for fraudulent transactions. It appears Target never spotted the attack until it was alerted by outside financial firms of a suspicious correlation between fraudulent charges and prior transactions at Target stores, according to security researcher Brian Krebs, who broke the story. Thus, it's likely that advanced exploits will be exposed only once the perpetrators start monetizing their pilfered information.

More than half of respondents to our security survey cite cyber-criminals as a top threat, with 37% very or extremely concerned about cyber-espionage by criminal syndicates or nation-states. We expect that number will be higher in our 2014 survey.

Core Principle: Abuse Internet Protocols To Mask Attacks

As we wrote in this column after viewing several convincing demos at a cyber-security conference, attacks using advanced evasion techniques are creative and crafted to exploit

Figure 4

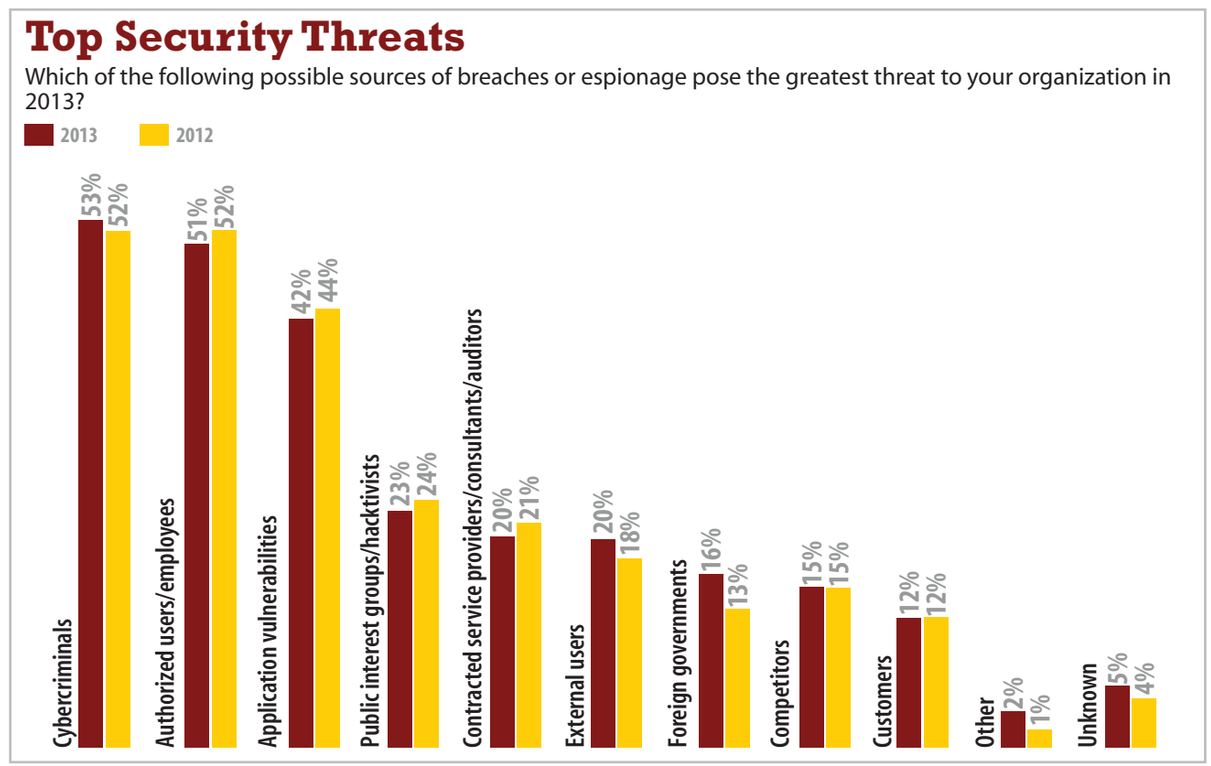

Note: Three responses allowed
Base: 1,029 respondents in March 2013 and 946 in March 2012
Data: InformationWeek Strategic Security Survey of business technology and security professionals at organizations with 100 or more employees

R6820513/13

weaknesses in the Internet's underlying technology. They take advantage of the inherent flexibility in how Internet protocols are written and implemented.

This is succinctly stated in the foundational IP IETF standard, RFC 760 (key points italicized). "The implementation of a protocol must be robust. Each implementation must expect to interoperate with others created by different individuals. While the goal of this specification is to be explicit about the protocol, there is the *possibility of differing interpretations.* In general, an *implementation should be conservative in its sending behavior, and liberal in its receiving behavior.* That is, it should be careful to send well-formed datagrams, but *should accept any datagram that it can interpret.*"

In other words, Internet participants are advised to drive by the rules of the road, but don't expect others to always do the same.

It is this asymmetry between sending and receiving behavior that advanced evasion techniques exploit. This is done by either manipulating IP-layer packets or TCP-layer data streams, altering low-level parameters, like TTL, packet length, and sequence numbers, or fragmenting streams in ways that look normal and safe to security appliances and operating systems but that can be reassembled into malware by malicious code on the host.

This is bad news for the security tools enterprises most rely on: 62% of respondents to our security survey cite firewalls as their most effective security practice, a higher rating than any other technique, while endpoint protection is mentioned by just 46%. The most devious and dangerous advanced evasion techniques are designed to bypass all of these.

Let's look at two particular categories.

Figure 5

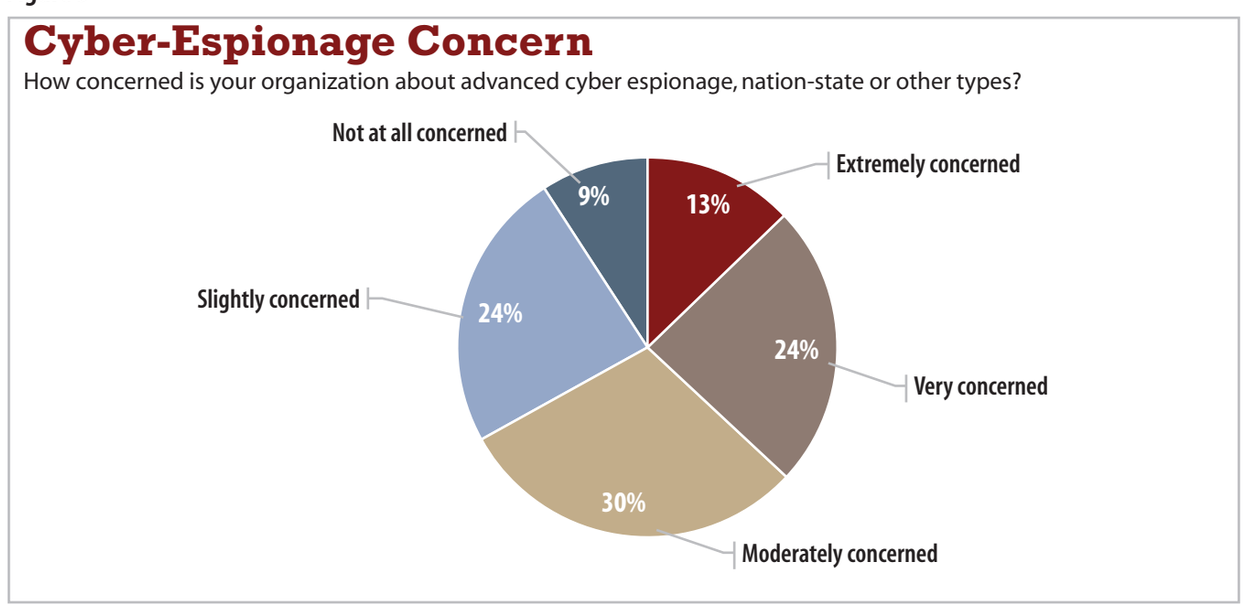

Data: InformationWeek 2013 Strategic Security Survey of 1,029 business technology and security professionals at organizations with 100 or more employees, March 2013

R6820513/14

1. IP Layer Evasions

At the IP layer, there are four primary means of bypassing network security: packet insertion, packet fragmentation or segmentation, data overlap, and presequence chaff. Of these, the first two are probably the most effective.

In a 1998 paper [PDF], one of the earliest descriptions of advanced evasion techniques, Thomas Ptacek (now at Matasano Security) and Timothy Newsham (now at iSEC Partners), describe the idea behind packet insertion. "An IDS can accept a packet that an end-system rejects. An IDS that does this makes the mistake of believing that the end-system has accepted and processed the packet when it actually hasn't. An attacker can exploit this condition by sending packets to an end-system that it will reject, but that the IDS will think are valid."

This extraneous data, for example, adding unusual IP options or header fields, can cause an IDS to miss a valid attack signature. That's because intrusion-detection

Figure 6

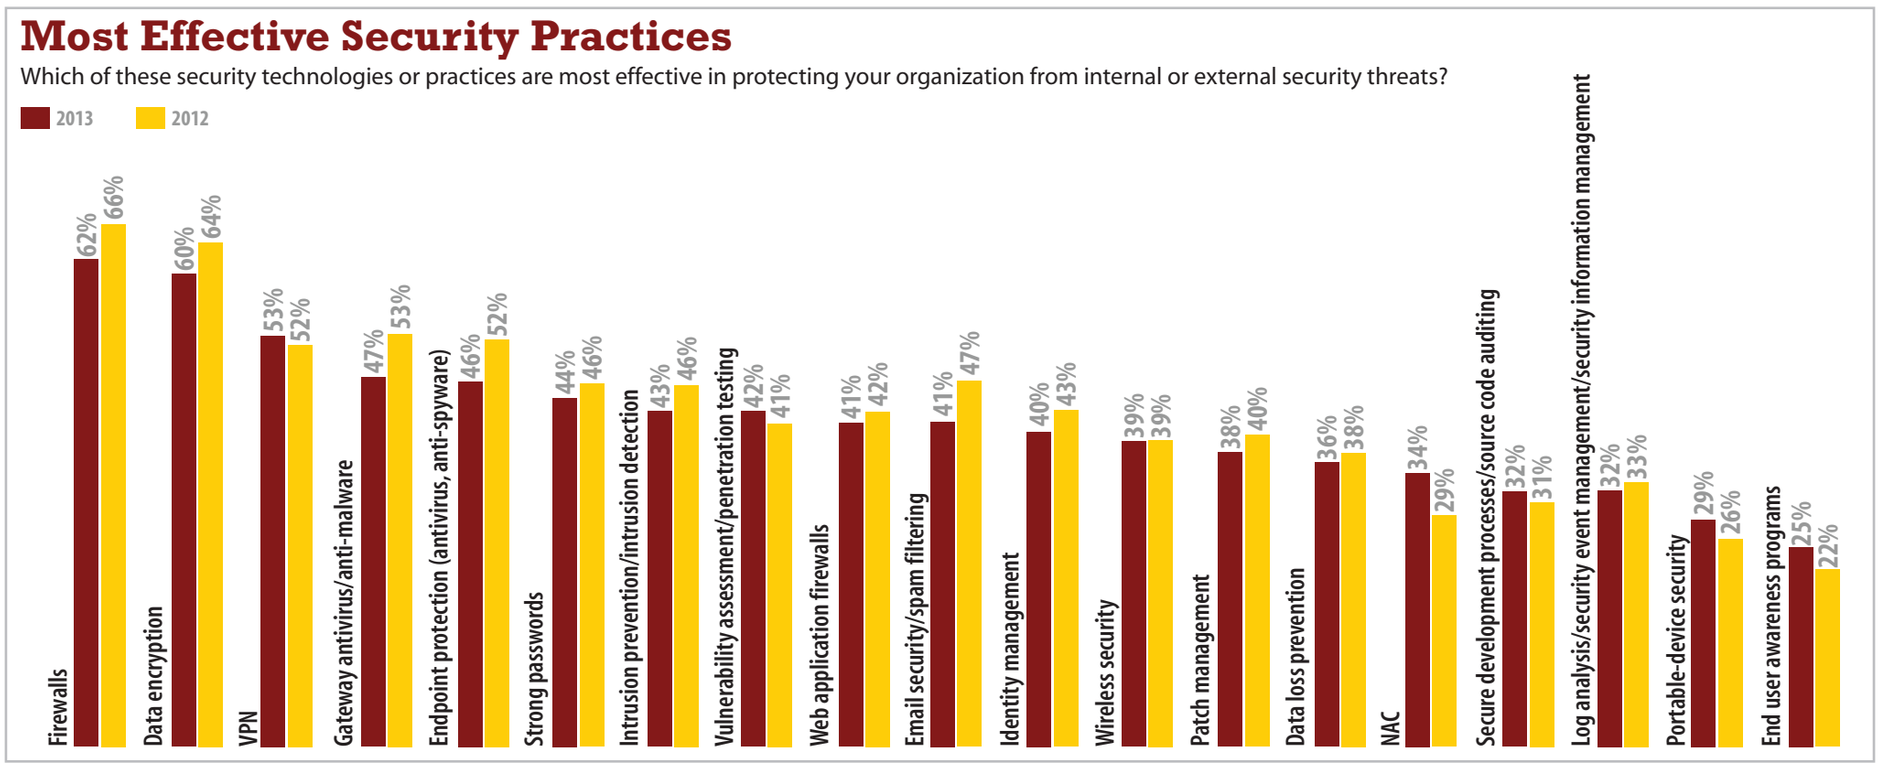

Note: Percentages reflect a response of "very effective"
Base: Respondents in March 2013 and March 2012 using each security technology or practice (varies)
Data: InformationWeek Strategic Security Survey of business technology and security professionals at organizations with 100 or more employees

R6820513/17

systems typically use a basic pattern-matching algorithm. Presequence chaff works similarly, by adding at least 1 byte of randomly generated data, prior to the data bytes in the packet header. The endpoint target will usually ignore this data, but an IDS may not and thus miss the fact that the data itself matches an attack signature.

Packet fragmentation or segmentation is an even more effective technique in that it exploits technical limitations of network security hardware itself. As a 2010 paper [PDF] by Stonesoft points out, in IP fragmentation evasions, the attacker takes advantage of scrambling fragments out of order, or by overwhelming the IPS with too many fragments. This allows malware to bypass IPS scanning by bits and pieces since fragments may be sufficiently separated in time that an IDS can't associate them as part of the same packet. It's analogous to a terrorist group smuggling a bomb through customs by sending fuses in one package, various chemicals in others. The result is a set of seemingly harmless items that can be combined to dangerous effect.

The inverse of fragmentation is data overlap, which, as the name implies, entails repeating the trailing bits of one packet as the beginning of the next. As Samuel Gorton and Terrence G. Champion, researchers at Skaion Corp., point out [PDF], since the data is correctly labeled, this is perfectly legal, and a network IDS may not correctly remove all duplicate data bytes. However, the duplicate data will also cause pattern-matching scanners to miss legitimate malware signatures.

2. TCP And Application-Layer Techniques

The most effective TCP-layer evasion, using out-of-order or overlapping segments, works much like IP fragmentation. Here, changing the TCP sequence order from a simple numerical increment to a random increment or decrement greatly complicates packet reassembly by the IDS and increases the amount of information it must buffer before completely processing the connection stream. Other TCP-layer techniques designed to create TCP stream reassembly problems and missed packets at the IDS include abusing TCP options, requiring three-way handshakes and exploiting faulty handshake detection, SYM packet abuse, packet overlap, mistaken TCP teardown, FIN and RST message processing, and even denial-of-service attacks designed to exhaust IDS resources.

At the application layer, evasions are typically aimed at the most common protocols: SMB (Windows file sharing), Microsoft RPC (remote code execution), and the big kahuna, HTTP. In the first two cases, the basic strategy remains the same — play a shell game with the IDS by mixing up traffic. On SMB, attackers can slice writes into multiple segments, perhaps even multiplexing them to different sockets (named pipes) over a single connection, and then reassembling them into malware on the host. Over MSRPC, a basic evasion involves switching the bit order of a connection. Since Windows typically uses little endian, an exploit sent as big endian may slip through a host's local defenses.

HTTP evasions are perhaps the most creative since the protocol is used as a conduit for increasingly flexible executable code

written in HTML5 and JavaScript. One possibility involves steganographically encoding images with embedded malware code in the bitstream of an image file. The end user and content filtering software just sees a picture. Decryption code embedded within JavaScript sees a valuable digital payload. As we wrote in our earlier Network Computing column, once downloaded, the code/image must be decrypted, extracted, and executed on the target. Here's where HTML5, with its rich JavaScript and CSS support, comes in. Essentially, a small piece of HTML5 code on the malicious web page can embed the necessary decryption instructions such that when the victim visits a website, an innocuous-looking image is automatically downloaded. Malware is extracted and then executed via an embedded shell script.

Bypassing a local machine's anti-malware defenses isn't particularly difficult for determined attackers. The usual technique involves obfuscating the payload using a polymorphic execution mangler/encoder to avoid a signature match. In fact, the popular Metasploit penetration testing tool includes a module (shikata_ga_nai) to do just this. Code obfuscation is even possible by wrapping the exploit in a virtual machine. That's an ironic twist since some security systems, like VMProtect, use this technique to create a safe sandbox for untrusted applications. However, attackers can use the same approach to execute arbitrary code and use polymorphically disguised entry and exit points to access underlying system resources without triggering an alert.

If you're thinking these techniques work only because they exploit decades-old technology and that once the world moves to IPv6 all will be well, think again. Some security researchers contend IPv6 will actually compound the problem because v6 software stacks are not as mature (and thus likely to have new bugs) while the protocol greatly expands the address space available to originate attacks, thus rendering blacklisting systems impotent. Robert Lemos explores IPv6 implications in more depth in a Dark Reading column.

The bottom line is that endpoints (targets) and network security scanners may not see advanced evasion techniques that couch harmful exploits as anything unusual.

Of course, since attackers can eventually untangle malware passed using advanced evasion techniques, why can't security appliances? They could, but at the cost of system resources and throughput. Security software needs to keep the entire attack stream in memory in order to detect it. But memory is expensive and can slow performance, particularly when the device is trying to handle millions of connections per section. As a simplified example, imagine that an IDS can hold 10 seconds of each packet stream in memory. If an attack is split into two fragments, 15 seconds apart, it will pass through undetected as just so much random packet noise. Even expanding memory buffers may not help since advanced techniques can be combined. For example, the executable might be wrapped in an encoded image sent via HTTP with the encryption key and the payload unpacker passed using packet fragmentation and extraneous headers. The number of possible permutations is in the billions.

FAST FACT

62%

of respondents to our security survey cite firewalls as their most effective security practice.

Are You At Risk?

While the situation seems dire, most enterprises needn't panic. Advanced evasion techniques are primarily the tools of determined cyber-criminals or nation-states and thus saved for APTs aimed at high-value targets like industrial control systems, utilities, and financial systems. Among security survey respondents reporting a breach in the past year, 25% say it involved theft of identity, IP, or confidential information, and 23% report financial losses. Overall, about 2% of our total sample of 1,029 respondents at organizations with 100 or more employees suffered losses of more than 5% of annual sales or revenue or couldn't estimate the damage. Ouch.

Although detecting advanced evasion techniques is difficult, it's not impossible. Next-generation firewalls from the likes of Cisco, Dell (SonicWall), McAfee (Stonesoft), and Palo Alto Networks include technology to detect attacks trying to leverage complexities in network and application protocols to bypass them. The basic approach is known as "traffic normalization," knitting packets containing various threads of an attack back to their intended state before inspecting for suspicious code or exploit signatures. However, IT teams must do their homework by understanding how vendors define advanced evasion techniques and pinning

Figure 7

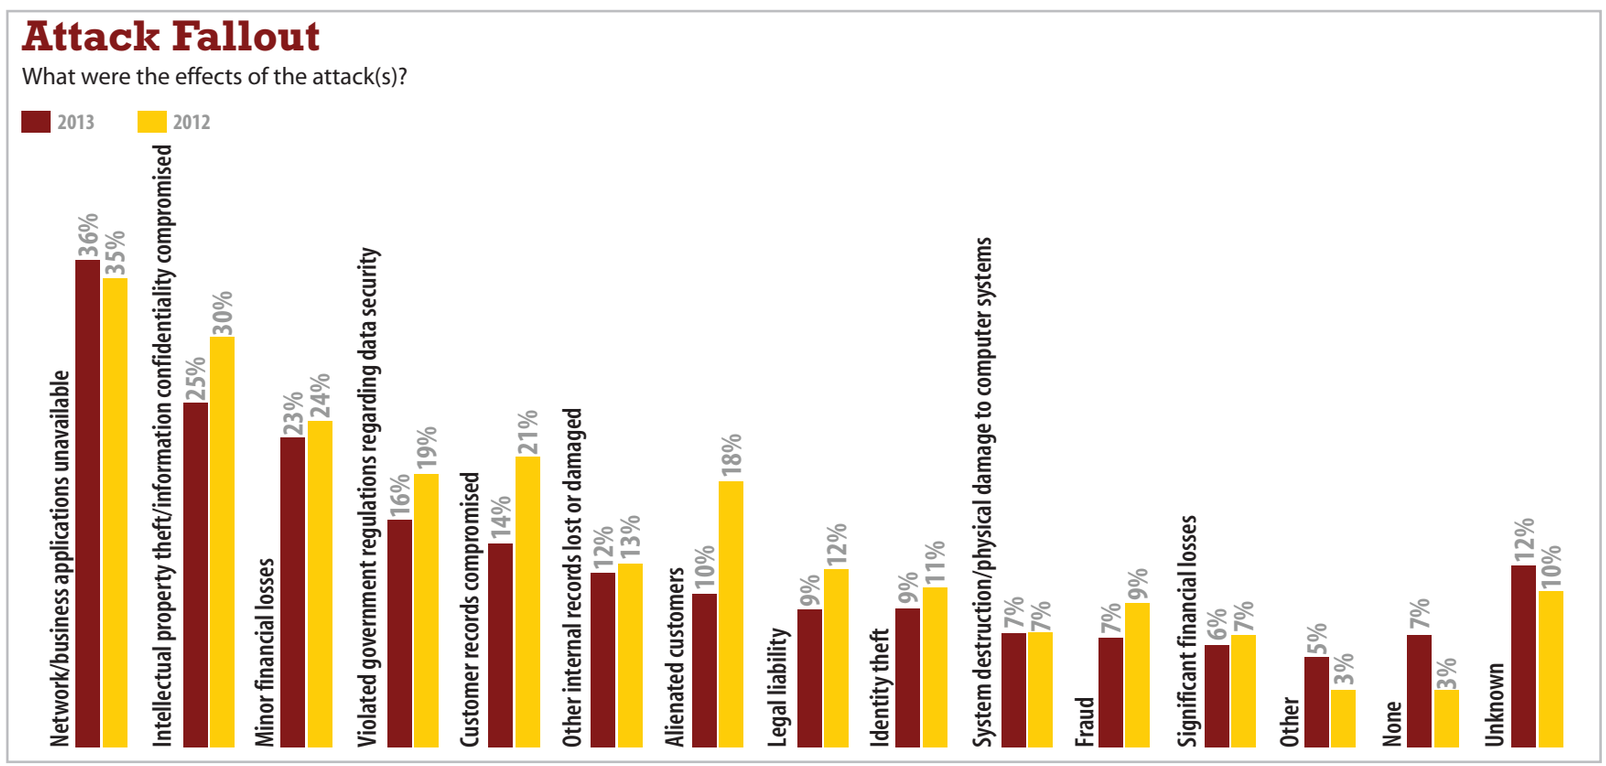

Note: Multiple responses allowed
Base: 217 respondents in March 2013 and 183 in March 2012 experiencing a security breach within the past year
Data: InformationWeek Strategic Security Survey of business technology and security professionals at organizations with 100 or more employees

them down on the specific evasions their software can detect. A poorly implemented product can give a misplaced sense of security by reporting false negatives, stating that an attack has been blocked when it actually slipped through.

The best way to evaluate a product's effectiveness is by testing it against actual attacks. Short of rolling your own Metasploit scripts, the quickest and easiest approach is to download Stonesoft's free Evader tool. Unlike Metasploit, Evader isn't a full-fledged pentesting suite, but it will unleash a variety of advanced evasion techniques to determine if a known exploit can break through your existing security devices and be delivered to target systems.

For organizations at high risk of loss, a budget for such detection should be a part of their security strategies going forward. Unfortunately, our data shows security spending, while increasing, is still tight. Of security survey respondents with knowledge of their organizations' spending plans, 70% devote 10% or less of their IT budgets to the task, although about one-third expect security spending to increase this year.

This isn't to say security isn't a priority; in fact, when we asked respondents to our latest IT Spending Priorities Survey to identify their top two initiatives for 2013, 35% named improving security, making it the most commonly listed project. However, when prodded to prioritize projects, improving security and regulatory compliance finished near the middle of the pack, at No. 4, behind improving business results, internal customer service, and operation of existing systems. Sadly, when push comes to shove, security often

FAST FACT

2%

of our 1,029 security survey respondents at organizations with 100 or more employees suffered losses of more than 5% of annual sales or revenue or couldn't estimate the damage.

Figure 8

Financial Losses

Approximately what percentage of your organization's annual sales or revenue did these financial losses represent?

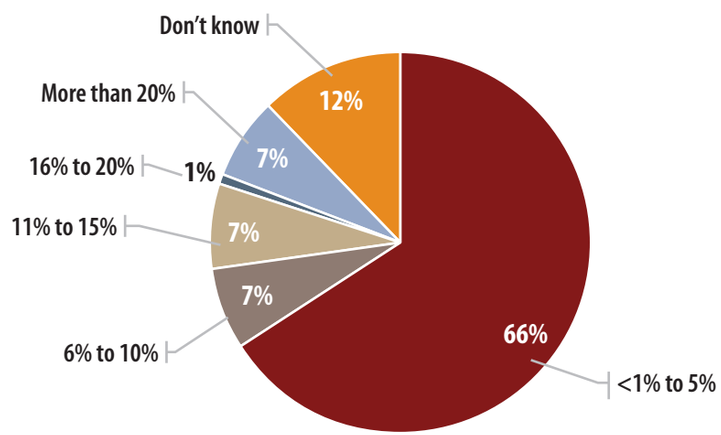

Base: 63 respondents at organizations experiencing significant or minor financial losses as a result of an attack
Data: InformationWeek 2013 Strategic Security Survey of 1,029 business technology and security professionals at organizations with 100 or more employees, March 2013

R6820513/7

takes a back seat to more pressing projects.

One of our respondents decried a general lack of interest in security by senior management: "It is difficult to get the executives and managers to appreciate the need for security at times and to understand that it's not as simple as installing an AV program and calling it done."

This has to change if we're ever to see a reduction in the types of massive, disruptive, and costly incidents epitomized by the Target breach. When it comes to the lengths attackers will take to circumvent network security, if the history of computer security demonstrates anything, it's how rapidly attack and evasion technology can disseminate. Prudence dictates lining up your defenses before advanced evasion techniques become standard fare in every attacker's toolkit.

Figure 9

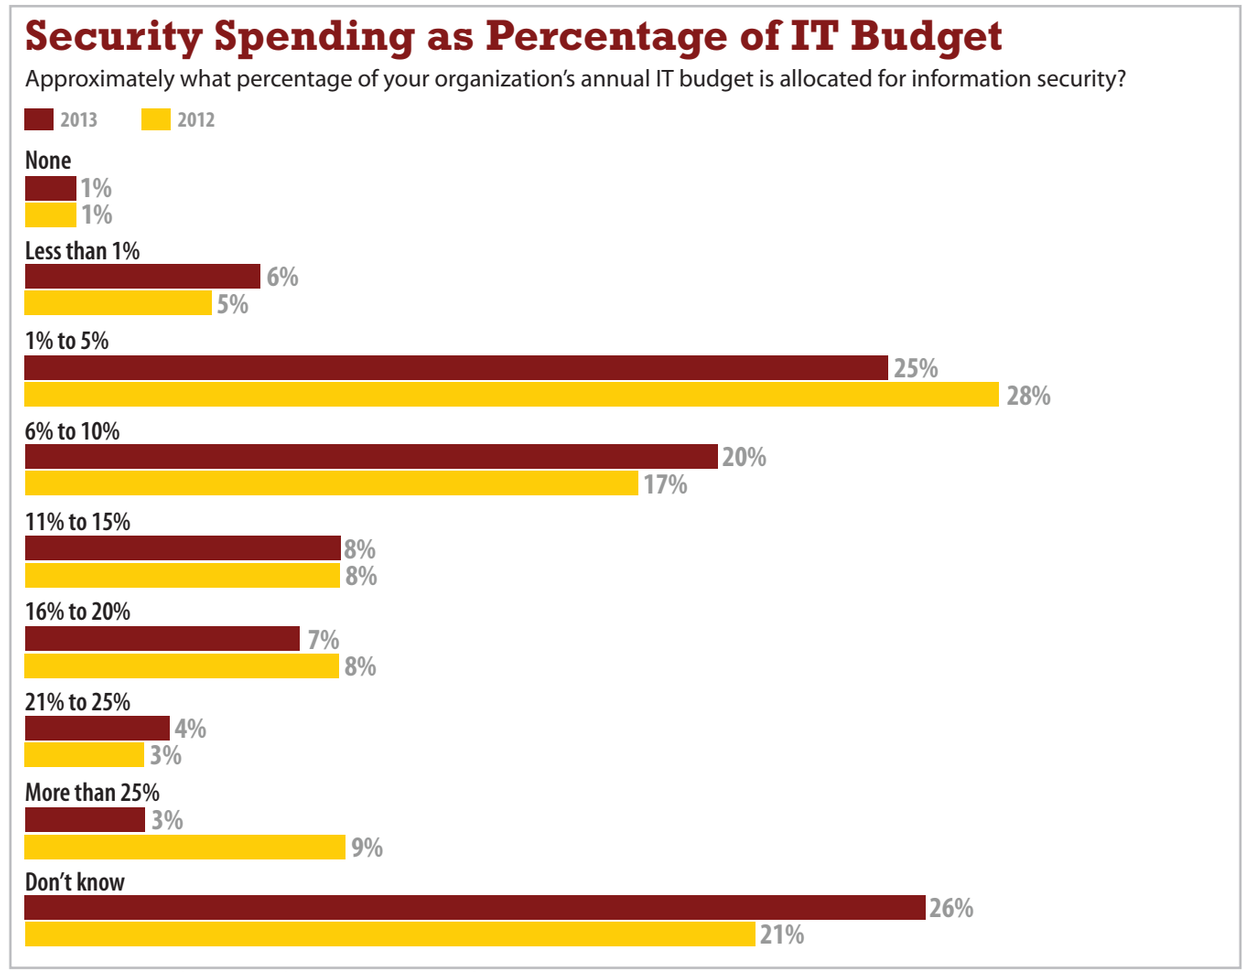

Base: 1,029 respondents in March 2013 and 946 in March 2012 R6820513/21
Data: InformationWeek Strategic Security Survey of business technology and security professionals at organizations with 100 or more employees

MORE LIKE THIS

Want More Like This?

InformationWeek creates more than 150 reports like this each year, and they're all free to registered users. We'll help you sort through vendor claims, justify IT projects and implement new systems by providing analysis and advice from IT professionals. Right now on our site you'll find:

The IPS Makeover: Next-gen intrusion-prevention systems have fuller visibility into applications and data. But do newer firewalls make IPS redundant?

Protecting Your Enterprise From DNS Threats: DNS is the world's largest distributed database, and it's increasingly being used as a launchpad for attacks. To gird against this growing threat, enterprise security pros need to understand the dangers posed and the damage that can be done. In this Dark Reading report, we provide a detailed examination of the ever-looming DNS threat, as well as advice and recommended resources for protecting the enterprise against it.

Creating and Maintaining a Custom Threat Profile: Intelligence feeds provide organizations with a wealth of data about security threats, but the information is valuable only if it is relevant and actionable. To be effective, intelligence data must be parsed and prioritized based on your company's specific IT systems and industry. In this Dark Reading report, we examine the types of information threat intelligence feeds can offer and recommend ways in which companies can most effectively develop a threat profile that will help keep their systems, data, and customers more secure.

PLUS: Find signature reports, such as the InformationWeek Salary Survey, InformationWeek 500 and the annual State of Security report; full issues; and much more.