

stop

badware

The State of Badware

June 2011

The logo features the word "stop" in a bold, red, sans-serif font, enclosed within a red rectangular stamp-like border. Below "stop", the word "badware" is written in a larger, dark blue, sans-serif font. The entire logo is centered at the top of the page.

StopBadware provides tools and information that assist industry and policymakers in meeting their responsibility to protect users from badware, and that help users protect themselves. It began as a project of the Berkman Center for Internet & Society at Harvard University before spinning off as a stand-alone nonprofit organization in 2010. Corporate partners include Google, PayPal, Mozilla, Nominum, Verizon, and Qualys. StopBadware is based in Cambridge, Massachusetts.

Visit us online at:

<http://www.stopbadware.org>

© 2011 by StopBadware, Inc.

All trademarks are the property of their respective owners.

This work is licensed under the Creative Commons Attribution-NoDerivs 3.0 Unported License.

To view a copy of this license, visit

<http://creativecommons.org/licenses/by-nd/3.0/>

Executive Summary

Badware — computer viruses, spyware, scareware, and the like — has emerged in the past decade as a key instrument in the perpetuation of cybercrime. While high-profile data breaches and targeted cyberattacks capture the headlines, it is the millions of infected home and business computers participating in botnets, stealing passwords, and scaring users into buying fake products that drive much of the underground economy. How do we understand the extent and evolution of this badware threat? What factors allow badware to succeed despite our best efforts to fight it? What is being done — and what should be done — to remove badware from its current position as a major driver of cybercrime?

Badware and its effects are notoriously difficult to measure, yet all signs point to continual growth in both the quantity and the complexity of threats. Each year shows a dramatic increase in the number of badware variants, complicating detection. Criminals move from attacking one application to the next, seeking out widespread and vulnerable targets. Websites, ad networks, and search engines increasingly serve as conduits for infecting unsuspecting users. Pools of millions of infected computers are now available for rent by malicious actors. The damages resulting from all this measure well into the billions of U.S. dollars.

How do we explain such a disturbing trend, even after years of efforts to combat badware and related cybercrime? The answer lies in four broad areas of vulnerability within the Internet ecosystem: technical, behavioral, economic, and legal. Today's software is often more secure than yesterday's, but challenges associated with patching applications and detecting ever-changing badware still abound. Internet users are faced with endless choices, but lack the knowledge or tools needed to elect the safest options. Market incentives drive businesses to make decisions that sacrifice security to other priorities. And our laws, policies, enforcement funding, and diplomatic agreements often predate the emergence of cybercrime as a pervasive threat — at the cost of investigating and prosecuting bad actors.

The challenge before us — to address these fundamental weaknesses in the ecosystem — is substantial, but not insurmountable. In the past couple years, several new initiatives have demonstrated the potential for public, private, and nonprofit players — often working collaboratively — to chip away at systemic flaws. Leading software vendors are beginning to change the dynamics of how applications are kept patched and how security software protects users from the latest threats. Those users are being assisted in their choices by more sensible defaults and through real-time warnings that prompt action and provide guidance when it is most needed. Businesses like ISPs and web hosting providers are facing incentives that make investments in security more compelling. Even legal systems are starting to catch up, as new policies and collaborations empower governments to better address domestic and international cybercrime.

These initiatives serve both as examples of what can be done and as reminders of how much is left to do. We — all of us who rely on and build upon the strength of the Internet as an engine of progress — must ensure that badware does not continue to undermine the trust necessary to sustain that progress. Stopping badware means a commitment not only to bolstering our own defenses, but to making major structural changes in how technology, behavior, economics, and policy interact in the ecosystem.

Introduction

The rise and persistence of badware are defining challenges for networked technology users, businesses, and governments across the world. Badware, in the broadest sense, is any code that fundamentally disregards users' choices about how their computers or network connections are used; it necessarily encompasses the entire range of software the security community refers to as malware, spyware, viruses, worms, Trojan horse programs, bots, and the like. Addressing the threats that badware poses requires stakeholders in the Internet ecosystem to understand not only what is known about badware's sources and prevalence, but also the structural factors that contribute to the current state of affairs. This report is intended as a resource that business leaders, policymakers, and concerned members of the public can use to understand the badware landscape, how it is evolving, and what concerned stakeholders are doing to address the threat.

In this report, we focus on the most common forms of badware: those that infect computers opportunistically with a financial motive. While we anticipate addressing other types of badware in future work, we exclude from this report analysis of targeted attacks (e.g., advanced persistent threats), use of badware in state-sponsored cyberwarfare, and new platforms (e.g., mobile devices) that have not yet matured as mainstream badware vectors.

We begin with an overview of the practice of measuring badware: we use current badware metrics to explore the various points at which security researchers are contributing to the collective understanding of the type and severity of the threats users face. Next, we identify a range of technical, behavioral, economic, and legal considerations that complicate efforts to reduce the spread and persistence of badware. Finally, we highlight and explore several examples of how the Internet ecosystem is responding to these challenges.

Measuring Badware

Although the negative consequences of badware — and thus the seriousness of the problem it poses — are obvious to any individual or organization that has been a badware victim, there are a number of considerations that limit the ways in which the security community can describe and examine the problem of badware as a whole. First, since the creators of modern badware actively seek to conceal badware's method of operation and presence on compromised computers, any individual piece of badware must first be detected by the security community before it can be counted towards infection (or disinfection) rates. Second, when a given piece of badware is detected by a researcher or reported to a particular security firm, it is generally assigned an identifier in accordance with the firm's policies. As a result, widely prevalent badware is often 'named' differently by a wide range of security vendors, making it difficult to compare and aggregate detected levels of a particular infection across the security landscape. Third, as the security community detects and classifies 'types' of badware, badware creators continually develop new ways to package and deploy badware in ways that defy existing classification methods.

Similar obstacles present themselves in measuring badware at points of distribution. When a security firm detects a web page that can cause its visitors to become infected with badware, it seems reasonable to count that page as a source of badware — but whether the unit of measurement is the URL of the individual page, the entire site associated with that URL, the domain name at the root of the URL, or the numerical IP address to which the URL resolves is a question answered differently across organizations.

Perhaps the most challenging aspect of badware to quantify is the damage it causes. Relevant metrics may include amounts of time and money spent mitigating attacks on infrastructure, measures of lost productivity, the value of stolen intellectual property, and resources expended by law enforcement on investigation, among others. In some cases, the impact or damage caused by badware infection may not be apparent and therefore may not be represented in damage estimates. Moreover, some forms of damage caused by badware may be difficult to quantify financially but are significant nonetheless, like loss of reputation.

Exacerbating these challenges is a lack of centralized reporting of badware data or incidents. One reason for this is the aforementioned challenge of common forms of classification and measurement. Another reason is the lack of incentives or requirements to report information that may — depending on the context — be viewed as trivial, embarrassing, or proprietary. Finally, the institutions to collect, organize, and effectively use the data do not yet exist.

With these caveats in mind, the metrics identified below explore different approaches to using publicly available information to understand the badware problem; they generally track the processes by which badware is developed, deployed, and employed by the criminal ecosystem responsible for it.

Samples

One way to envision the evolution of badware is by examining the number of malicious code samples the security community has collected. A sample is a unique application or other set of code; duplicating this code results in an identical sample, while modifying the code results in a new sample. A-V Test Labs collected just over 17.5 million unique samples of malware code in 2010, up from fewer than 12.4 million in 2009, a 41 percent increase.¹ Similarly, Panda Security’s antivirus products detected approximately 59 million unique malware variants in 2010, up from approximately 38 million in 2009 — a 64 percent increase.²

New Malware Samples added to AV-Test Labs Repository




Figure 1. A rapid increase in the number of unique malware samples makes timely detection an ongoing challenge for security vendors.

Data reproduced with permission from AV-Test.org (<http://www.av-test.org>).

This mode of measurement suggests a substantial increase in efforts by malware authors to hide their work from security software. Badware detection methods that rely on identifying specific badware samples (so-called signature-based detection) can miss the spread of badware that is dynamically altered to create numerous new samples with identical behavior. Furthermore, since badware source code can be (and is) repurposed and repackaged by the criminal syndicates backing badware authors, growth in the number of unique badware samples collected ‘in the wild’ may indicate a greater number of active criminal syndicates using badware. Despite media reports using the growth in sample numbers to proclaim “increases in malware,” it is not necessarily the case that an increasing diversity of samples correlates to an increase in user exposure or a rise in successful infections.

¹ “Year-end malware stats from AV-Test.” GFI Labs Blog: January 27, 2011. Available at <http://sunbeltblog.blogspot.com/2011/01/updated-virus-stats-from-av-test.html>.

² “The Cyber-Crime Black Market: Uncovered.” Panda Security: January 20, 2011. Available at <http://press.pandasecurity.com/wp-content/uploads/2011/01/The-Cyber-Crime-Black-Market.pdf>.

Vulnerabilities

Another approach to measuring badware focuses on the types of software vulnerabilities that badware families are exploiting to gain unauthorized access to users' computers. Since badware authors often aim to maximize the impact of each exposure to malicious code, badware is generally likely to target weaknesses in widely installed software that users trust to accomplish daily tasks. In 2010, the security community observed a substantial increase in the exploitation of three nearly universally installed applications. Vulnerabilities in Oracle's Java Virtual Machine came under heavy attack, particularly in the latter half of the year. As illustrated in Figure 2 below, Microsoft desktop anti-malware products detected attempts by badware to exploit Java over 6.5 million times in the third quarter of 2010, and over 6 million attempts in the fourth quarter, up from under 500,000 attempts earlier in the year. By the year's end, Java exploits made up nearly half of all exploit attempts Microsoft was able to detect.³ Kaspersky Labs discovered that one particular exploit method for Java rose from complete disuse in October 2010 to approximately 40,000 unique infections per day at the end of January 2011; three Trojan horse programs relying on the exploit accounted for over 260,000 infections detected by Kaspersky in January 2011 alone.⁴




Figure 2. Exploit attempts by technology targeted, as detected by Microsoft in 2010. Criminals have increasingly attacked popular cross-platform applications such as Java. Graph reproduced with permission from Microsoft Security Intelligence Report, Vol. 10.

Adobe Reader, which millions of computer users rely on to view PDF files, also came under substantial attack. According to GFI Labs, Adobe PDF vulnerabilities accounted for 2 out of the top 10 most detected malware threats in December 2010.⁵ Kaspersky Labs reported that the most prevalent badware packaging utilities ('exploit kits') used Adobe Reader as an attack vector in 28 percent of cases.⁶ Similarly, Adobe Flash, which enjoys a nearly universal base of installation due to Flash's popularity as a rendering engine for streaming video and dynamic websites, was targeted by the Phoenix badware packaging utility in 20 percent of cases.⁷

³ Cavit et al. Microsoft Security Intelligence Report, Vol. 10: 20. Microsoft: May 12, 2011. Available at <http://www.microsoft.com/security/sir/default.aspx>.

⁴ Zhakorzhovsky, Vyacheslav. "Monthly Malware Statistics, January 2011". Kaspersky Labs: February 3, 2011. Available at http://www.securelist.com/en/analysis/204792159/Monthly_Malware_Statistics_January_2011.

⁵ "GFI's Top 10 Malware List." GFI Labs: February 4, 2011. Available at <http://www.gfi.com/page/67883/january-sees-uptick-in-targeted-attacks-on-adobe-reader-files-according-to-gfis-top-10-malware-list>.

⁶ Diaz et al. "Fuel for pwnage: Exploit kits." Presentation at SOURCE Boston Conference: April 11, 2011. Available at <http://www.slideshare.net/SOURCEConference/vicente-diaz-jorge-mieres-fuel-for-pwnage-777711>.

⁷ *Ibid.*

Infection Vectors

Focusing on software vulnerabilities gives an indication of the risk any given computer may face when exposed to a given piece of badware, but how likely is it that users will encounter that badware in the first place? Measuring vectors for badware infection is essential to understanding the badware threat. Infection vectors are substantially more difficult to quantify than vulnerabilities because some badware (worms) can spread from computer to computer without any user intervention; furthermore, users and computers may manipulate badware-infected files in many ways — for example, a maliciously crafted PDF might be received via e-mail and opened by an unsuspecting user, or downloaded from a flash drive, or loaded by visiting a website.

Visiting websites is one of the most common ways users are exposed to badware. Some sites are purpose-built to cause computers to install badware, but many are ordinarily legitimate sites that have been compromised to serve code that redirects the user to badware content without any obvious indication that such content is being loaded. The threat that ‘badware websites’ pose appeared to be on the rise in 2010. Websense reported a 111.4% increase in the number of malicious websites it observed over its 2009 figures, without disclosing the raw number of sites used to calculate the increase.⁸ Dasient, a security firm specializing in the detection and remediation of badware websites, observed 1.1 million infected websites at the close of 2010, up from approximately 550,000 such sites at the close of 2009.⁹ StopBadware’s own analysis of URLs reported as badware by its data providers — Google, GFI Labs, and NSFOCUS — shows a similar increase in activity: at the beginning of January 2010, about 350,000 URLs were reported; this number rose to over 575,000 by the year’s end. Early figures for 2011 show a continuing upward trend.

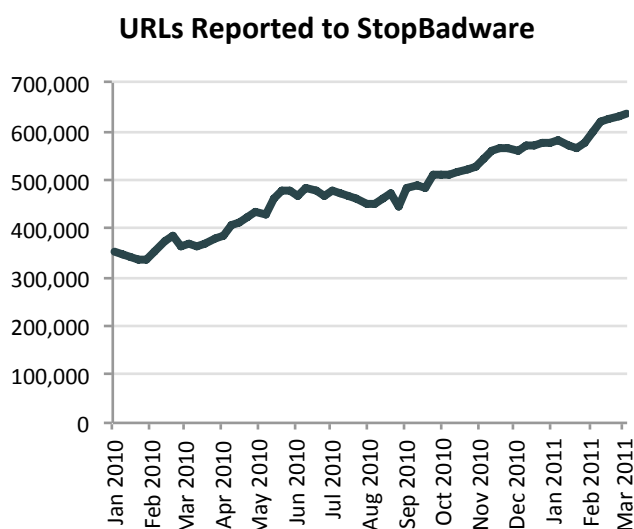


Figure 3. Badware websites have become an increasingly popular vector for badware delivery, and the threat continues to grow.

URL data provided by Google, GFI Labs, and NSFOCUS.

While such measurements indicate an increase in detected badware websites, they do not address how likely users are to visit those sites. One of the most common ways users find content on the Internet is through search engine results. Both Google and Microsoft, whose Google Search and Bing Search products captured just over 95 percent of the U.S. search engine market at the end of 2010,¹⁰ have attempted to protect their users from malicious web sites — Google through its Safe Browsing initiative, and Microsoft through its SmartScreen filter. Both companies have reported fairly low levels of exposure to drive-by downloads in search results: in July 2009,

⁸ Websense 2010 Threat Report. Websense: November 9, 2010. Available at <https://www.websense.com/assets/reports/report-websense-2010-threat-report-en.pdf>.

⁹ “The Dasient Q4 Malware Update.” Dasient: March 7, 2011. Available at <http://blog.dasient.com/2011/03/dasient-q4-malware-update-significant.html>.

¹⁰ “Bing Continues Growth in January 2011.” Search Engine Watch: February 22, 2011. Available at <http://searchenginewatch.com/article/2066029/comScore-Bing-Continues-Growth-in-January-2011>.

Google indicated that approximately 0.75 percent of all Google search queries returned at least one site flagged as malicious.¹¹ In mid-2010, Microsoft indicated that drive-by download pages accounted for 0.3 of every 1,000 pages in the Bing index and appeared on 2 out of every 1,000 search pages displayed to users.¹² These metrics may substantially underestimate the level of risk to which users are exposed, however. First, sites that have been compromised after a search engine's most recent visit will remain undetected — and users will remain unwarned — until such sites are appropriately identified. Second, cybercriminals actively manipulate the content of some badware websites to increase such sites' search ranking for timely or popular search queries. Websense reported that in 2010, 22.4 percent of popular trending terms on Google and Yahoo search engines linked to malware, as opposed to 21.8 percent for known sex terms.¹³ Instructing users to avoid 'risky' (or here, risqué) content in searches, therefore, is not enough to reduce users' exposure to the threat of badware websites.

Base measurements of badware websites may also under-report the likelihood of users encountering badware due to a rise in malvertising. Web-based advertising, like web search, is a ubiquitous way in which users are exposed to new content; cybercriminals can expose large numbers of users to badware by exploiting the syndication model of ad networks and evading the networks' vetting practices. Any visitor to a site displaying such a malvertisement is at risk of badware infection until the ad is taken down. For example, visitors to londonstockexchange.com and several other high profile UK sites were exposed to fake AV in February 2011, via a malicious ad served by ad network Unanimis.¹⁴ The threat such malvertisements pose appears to be increasing. Dasient estimated that 3 million malvertising impressions were served per day in Q4 2010, up from 1.5 million impressions in Q3 2010.¹⁵

Endpoints

To further understand the scale and scope of the badware problem, it is desirable to measure the prevalence of badware after the point at which users' computers have become infected. Measuring numbers of infected computers faces all of the observational and monitoring limitations inherent in analyzing badware itself: the badware (in this case, the installed or running badware) must be detected, categorized, and reported. This requires computers, or their network traffic, to be monitored by tools that can detect the badware present on the systems. We will refer to infected devices of all types as 'endpoints' of the badware ecosystem.

Estimates of the number of infected badware endpoints may range widely based on the capabilities and orientation of the observing organizations. Panda Labs, using as its baseline a set of over 18 million computers running its malware scanning software, reported that over 9 million computers, or 50.3% of the total, were compromised by active or latent malware in the first half of 2010.¹⁶ Such a startlingly high percentage may be explained in part by the fact that users who have elected to submit their computers to malware scanning software are more likely than an average user to have observed undesirable badware-related behavior — self-selection, in effect.

Another proxy for the number of badware endpoints is the number of computers detected to participate in botnets. Bots are a class of badware specially designed to receive and execute instructions sent via malicious "command and control" computers run by cybercriminals; this converts infected computers into distributed platforms for negative behavior, including attacking websites to interfere with their use (known as "denial of service" attacks), sending spam (junk email) mailings, capturing user account credentials, and the like. While the

¹¹ Provos, Niels. "Malware Statistics Update." Google: August 25, 2009. Available at <http://googleonlinesecurity.blogspot.com/2009/08/malware-statistics-update.html>.

¹² Anselmi et al. Microsoft Security Intelligence Report, Vol. 9: 100. Available at <http://www.microsoft.com/security/sir/archive/default.aspx>.

¹³ Websense 2010 Threat Report (see note 8).

¹⁴ Leyden, John. "Tainted ads punt scareware to surfers on LSE and Myvue sites." The Register: February 28, 2011. Available at http://www.theregister.co.uk/2011/02/28/tainted_ads_blight_uk_sites/.

¹⁵ The Dasient Q4 Malware Update (see note 9).

¹⁶ Phishing Activity Trends Report: 2nd Quarter 2010. Anti-Phishing Working Group: January 26, 2011. Available at http://www.antiphishing.org/reports/apwg_report_q2_2010.pdf.

size of a botnet may not indicate its direct threat potential, the size of such networks can help us understand the secondary damage to endpoints.¹⁷ Botnets are notoriously difficult to measure, but there are some indicators available of their pervasiveness.

In April 2011, the United States Federal Bureau of Investigation estimated the Coreflood botnet to consist of “hundreds of thousands” of infected PCs.¹⁸ Damballa indicates that botnets of one hundred thousand to two million PCs are readily available for lease or purchase by malicious actors through underground marketplaces.¹⁹ Measurements gathered by Microsoft may be able to shed broader light on the number of badware endpoints; this is due in large part to the fact that Microsoft’s Malicious Software Removal Tool (MSRT) — the source of the disinfection measurements — is run monthly through Windows Update and is available for on-demand execution. It is therefore widely used and may gather information from users who have not actively pursued antivirus solutions. Its botnet activity measurements from 2010 are notable both for the high raw number of botnet badware disinfections occurring, and for the characteristics of the botnet badware being disinfected.

The top 5 botnet badware families detected and cleaned by Microsoft’s Malicious Software Removal Tool — Rimecud, Aleuron, Hamweq, Pushbot, and IRCbot, respectively — made up about 10 million computer cleanings in the first half of 2010.²⁰ In a worrisome development, moreover, one of these families, Hamweq, was observed to install Rimecud, a different family, leading to Rimecud being detected on Hamweq-compromised systems in 34.1% of cleanings.²¹ This observed consolidation is the more worrisome because all of the top detected botnet badware families are capable of downloading and installing other forms of badware on command, including updates to the botnet badware itself, allowing cybercriminals to further compromise and commoditize the resources of the infected machines. And Rimecud itself maintained a high and constant prevalence on compromised machines throughout 2010, varying between 1.67 million and 1.82 million detections per quarter, and was cumulatively detected at over 7.1 million MSRT executions in 2010.²² It’s notable that the number of cleanings may not be a perfect proxy for the number of currently infected endpoints: any single computer may have been cleaned multiple times, and a cleaned computer is no longer infected. Still, when one takes into account that botnet disinfections constituted a mere 33 percent of malicious software removed by the tool, and that the tool only removes selected families of badware,²³ it becomes clear that the overall number of active badware endpoints must be quite high.

Damage

Ultimately, from the perspective of individuals, businesses, governments, and other Internet stakeholders, one of the most important badware measurements is the cost imposed on badware victims. Some costs are obvious and easily monetized: for example, in October 2010, the FBI, working with international law enforcement, managed to shut down a group of cybercriminals using popular badware (known as ZeuS) to extract banking information from unwitting victims and conduct mass debits from their bank accounts; while the group was foiled in its attempts to steal \$220 million from their victims, they succeeded in transferring \$70 million out of the stolen accounts before they were caught.²⁴

¹⁷ Plohmann et al. “Botnets: Measurement, Detection, Disinfection, and Defense.” ENISA: March 7, 2011. Available at <http://www.enisa.europa.eu/act/res/botnets/botnets-measurement-detection-disinfection-and-defence>.

¹⁸ Zetter, Kim. “With Court Order, FBI Hijacks ‘Coreflood’ Botnet, Sends Kill Signal.” Wired: April 13, 2011. Available at <http://www.wired.com/threatlevel/2011/04/coreflood/>.

¹⁹ Ollmann, Gunter. “Building Botnets for Fun & Profit.” Damballa: October 9, 2010. Available at http://www.damballa.com/downloads/r_pubs/HackerHalted2010.pdf.

²⁰ Microsoft Security Intelligence Report, vol. 9 (see note 12).

²¹ *Ibid.*

²² Microsoft Security Intelligence Report, vol. 10: 41 (see note 3).

²³ “The Botnet Superhighway.” Microsoft: October 21, 2010. Available at <http://blogs.technet.com/b/mmpc/archive/2010/10/21/the-botnet-superhighway.aspx>.

²⁴ “Global Law Enforcement Cooperation Key in Disruption of Cybercrime Ring.” SecurityWeek News: October 4, 2010. Available at <http://www.securityweek.com/global-law-enforcement-cooperation-key-disruption-cybercrime-ring/>.

Other costs are harder to estimate but no less serious: in February 2011, the U.K. government estimated that cybercrime costs the British economy £27 billion per year — 2 percent of the country's GDP — and that three quarters of this cost was borne by industry.²⁵

Such costs are not confined to industry, of course. In a survey of 82 million American households, Consumer Reports found that 8 million households encountered serious problems with spyware in the first half of 2010; 617,000 households, presumably uncertain as to how to rid themselves of badware, indicated they had to replace slow or impaired computers, causing overall damage to consumers of \$1.2 billion.²⁶ Certain types of badware compromise, however, cause damage with no immediately quantifiable cost, instead causing a loss of reputation and organizational security. In March 2011, the French government confirmed that computers at its Ministry of Finance were compromised by e-mailed Trojan horse programs, causing the loss of sensitive documents related to France's G20 presidency.²⁷ Though the fact that such an attack was clearly targeted distinguishes it from much of the badware typology we have identified, the technique borrows heavily from the methods used for mass badware distribution, and the attackers likely benefited from the underground economy mass badware perpetuates.

Why are things in this state?

The above measurements afford security community stakeholders partial views into the broad contours of the badware landscape. Despite the fact that making rigorous claims based on this data is at best an exercise in speculation, certain trends are clear:

- The complexity of detecting and categorizing badware is increasing at a rapid rate;
- Badware authors are targeting exploits in widely installed, cross-platform applications;
- Badware is increasingly distributed through visits to websites serving drive-by downloads and malvertisements;
- Badware is compromising user data and commoditizing computer resources in increasingly granular ways; and
- Cybercriminals responsible for the development and distribution of badware have refined the self-reinforcing cycle of infections to establish a mature badware economy.

What accounts for this state of affairs? More pointedly, what structural factors are involved in facilitating and perpetuating badware infection? By examining the technical, behavioral, economic, and legal constraints that modulate the Internet ecosystem, we can begin to understand why it has been so difficult for the anti-badware community — victims, security companies, businesses, and governments alike — to systematically address the badware problem.

Technical factors

The infection surface any given computer presents for badware infection is meaningfully increased by the presence of unpatched, exploited vulnerabilities in the software that computer runs. Patching those vulnerabilities in a timely fashion is a difficult task, especially for large producers of widely installed and used software. Once a user installs a piece of software on a computer, the ability of that software's author to patch vulnerabilities is contingent upon either an automatic update mechanism or specific action on the user's part. Older software, in particular, is unlikely to default to the use of automatic updates.

²⁵ Bryan-Low, Cassell. "Cybercrime Costs Mount in U.K." The Wall Street Journal: February 17, 2011. Available at <http://online.wsj.com/article/SB10001424052748703561604576150353058208060.html>.

²⁶ "State of the Net 2010." Consumer Reports: June 2010. Available at <http://www.consumerreports.org/cro/magazine-archive/2010/june/electronics-computers/social-insecurity/state-of-the-net-2010/index.htm>.

²⁷ Cluley, Graham. "French Ministry hit by hacker attack, targeting secret G20 plans." Sophos: March 7, 2011. Available at <http://nakedsecurity.sophos.com/2011/03/07/french-ministry-hacker-attack-secret-g20-plans/>.

Further complicating defensive efforts, badware families with unique individual payloads are increasingly prevalent, forcing trade-offs for both users and antivirus companies. To ensure the highest degree of protection possible, antivirus products must regularly update their libraries of signatures and expose incoming files and network communications to a computationally intense degree of scrutiny; users must decide the degree to which they are willing to bear the delays such security measures impose.

Behavioral factors

While some badware infections can propagate without any form of user intervention (that is, through automatically exploiting vulnerabilities in software), badware often requires some sort of proximate user-computer interaction to infect a given computer. In some cases, users may be unaware that clicking a link or opening a file purporting to come from a trusted source can lead to infection; in others, users may treat badware infection as an annoyance to be dealt with rather than a threat to their (or their company's) data and computing resources.²⁸ These attitudes matter when users make decisions about when (or whether) to update their operating systems and other software. Faced with the task of keeping multiple programs up to date, users must balance the uncertain and downplayed risk of badware infection as a result of not updating with the cost of running updaters, downloading updates, and allowing those updates to install. Updaters that require higher amounts of user attention and intervention to patch software are therefore more likely to cause users to postpone updating — or ignore updates altogether.

The one factor most likely to prompt rapid user action — visible signs of badware, such as scareware popups — are not always present. Indeed, badware applications with such behaviors are unlikely to pose the greatest risk to a computer's integrity *because they call attention to themselves*. Trojan horse programs and botnet badware, in contrast, tend not to obviously advertise their presence.




Figure 4. Web pages that mimic the Windows user experience and deliver fake antivirus software are among the most visible threats to computer users.

²⁸ Wash, Rick. "Folk Models of Home Computer Security." Proceedings of the Sixth Symposium on Usable Privacy and Security (SOUPS '10). Available at <http://www.rickwash.com/papers/rwash-homesec-soups10-final.pdf>.

Economic factors

Addressing the problem of badware at a systemic level is a costly proposition and generally involves a large number of stakeholders, each of which may have only a limited ability to observe or manipulate the links in the Internet ecosystem that badware uses to propagate. Most stakeholders, moreover, are businesses or other service providers whose primary focus is not computer security; thus the amount of time, money, and other resources they spend combating badware will be determined by the degree to which badware threatens their own areas of interest. For example, commercial web hosting providers compete for customers based on the robustness of hosting features they offer, for a price; they have substantially less of an interest in controlling or monitoring the security of the content they host, since such a feature is difficult to systemically guarantee. When an individual customer's site is compromised by badware distributors, identifying and redressing that compromise may be one of several issues competing for attention and allocation of funds. A similar situation faces domain name registrars, who compete for customers on the basis of price, ease, and speed of domain name registration and management; when criminals register domain names for the explicit purpose of facilitating badware distribution, registrars have no financial incentive to investigate or revoke such registrations, or to take steps to proactively detect abusive registrations. In essence, badware operates in an environment wherein business efficiency incentives directly counter badware prevention and remediation procedures — procedures that are frequently costly and investigation-intensive.

To the extent that stakeholders do see economic value in assisting with the fight against badware, they must know what to do to help. A lack of best practices and uniform policies for identifying, reporting, and tracking badware incidents may hamper stakeholder efforts to contribute to the fight, particularly in the absence of available security expertise. Sharing data (such as logs indicating malicious activity, code from compromised webpages, strategies attackers use to compromise sites, or badware-infected files) is a time-consuming and frequently time-sensitive effort which can excite concerns about customer privacy and reputation. Without incentives to act and indemnity from the risks of acting in good faith, the rational economic response of stakeholders will be not to act. While the security community — antivirus companies, threat monitoring networks, vulnerability researchers, and others — are to a degree united in the goal of reducing the threat of badware, their approaches and areas of attention differ, based in large part on the needs of their customers. Sharing the latest badware data across the community in a way that balances for-profit companies' interest in economic competitiveness with the potential benefit of collaboration also poses a challenge.

Legal factors

Systematically addressing the problem of badware poses a number of important legal questions. It is beyond dispute that badware directly enables activities that are almost universally illegal, such as bank and credit card fraud, as well as those that some governments have chosen to criminalize, like sending mass unsolicited e-mail (i.e., spam). Targeting badware production, distribution, or operation explicitly as legal wrongs in themselves, however, raises important public policy questions. Since many stakeholders in the Internet ecosystem facilitate badware infection unwittingly (such as compromised site owners, web hosting providers, DNS providers, domain registries, and the like), the extent to which they can reasonably be held accountable is an open question. Moreover, any regulation of computers and the Internet runs into concerns about free speech, privacy, and users' freedom to engage in legitimate business online. Thus, such regulation remains in its infancy across the world.

Even when laws that can be used to punish badware distributors exist (like the United States's Computer Fraud and Abuse Act, 18 U.S.C. § 1030), law enforcement personnel require a substantial quantity of technical and financial resources to conduct investigations, coordinate with other participating agencies, bring charges against those responsible, and (if feasible) repair the damage badware has caused. The difficulty in overcoming these resource challenges has likely contributed to a lack of systemic prosecutions, and very little case law exists that specifically addresses the challenges law enforcement may face in constructing and bringing cases. Most pressing of all, because the architecture of the Internet transcends national boundaries, investigating and prosecuting badware distributors frequently requires the cooperation and assistance of courts and law enforcement personnel outside the investigating government's jurisdiction. Without well-defined channels for cooperation between governments, attempts to investigate badware problems may, and often do, stall.

Ecosystem Responses

As badware distributors have refined the tactics they use to compromise computers, other stakeholders within the Internet ecosystem have taken steps to address the structural factors that contribute to badware's prevalence and persistence. Security companies are developing increasingly dynamic and centralized antivirus technologies that more accurately target the proximate sources of computer infection, and large software distributors are more aggressively asserting the importance of software updates to reduce the prevalence of exploited vulnerabilities. National governments and private parties have embarked on a number of initiatives designed to detect and disinfect computers compromised by botnet badware, and to target the command and control networks that control the bots themselves. Independent organizations are further seeking to decrease the cost and increase the consistency of data sharing and badware response within the broader security community.

Software company efforts

Antivirus products are most users' last line of defense against badware. As badware distributors deploy increasingly sophisticated, signature-undetectable, and time-sensitive attacks on users, antivirus companies have altered their software offerings in several important ways. First, antivirus products have increasingly employed so-called cloud-based detection technologies when parsing incoming content for badware. Such products frequently update virus definitions and heuristics in real time from security vendor servers, and in some cases offload testing of files to analysis servers, reducing the performance impact of behavioral analysis on the user experience. These features are available in commonly used paid antivirus products (such as Symantec's²⁹) as well as in free solutions (such as Panda Labs'³⁰ and AVG's³¹). Second, antivirus products increasingly use reputational data and malware blacklists to protect users from web-based badware delivery mechanisms. Approaches differ: while AVG's LinkScanner purports to check linked web pages "in real time" before users click links,³² McAfee's SiteAdvisor cites its use of web crawlers to detect malicious site behavior.³³ Overall, users running such antivirus products are more likely than ever to have up-to-date protection information; when protection methods fail, the data such antivirus products now gather and transmit to the cloud can be used to improve detection.

Outside the antivirus community, developers of popular software have increasingly applied more sensible security defaults to their products. In 2010, Adobe, whose Flash Player and Reader applications are among the most frequently targeted for exploitation, took steps on its own and in collaboration with others to streamline update functionality. In April 2010, Adobe announced that its Reader 8 and 9 products would be configured to automatically download and install updates by default, rather than downloading updates and waiting for user intervention to install them; users who had modified the previously existing default settings were unaffected.³⁴ In May 2010, Adobe further announced plans to integrate updates to Adobe Flash into Google Chrome's internal update process.³⁵ By working to reduce the number of software update paths and the intrusiveness of update processes, software companies can improve the integrity of their install base and reduce user exposure to widely exploited software vulnerabilities.

²⁹ "Endpoint Protection.cloud." Symantec. Available at http://www.message-labs.com/products/hosted_endpoint/.

³⁰ "Download Panda Cloud Antivirus." Panda Security. Available at <http://www.cloudantivirus.com/en/download/cloud-antivirus/free/>.

³¹ "Why AVG?" AVG. Available at <http://www.avg.com/us-en/why-avg/>.

³² *Ibid.*

³³ "McAfee Site Advisor Enterprise." McAfee. Available at <http://www.mcafee.com/us/products/siteadvisor-enterprise.aspx>.

³⁴ Gottwals, Steve. "Upcoming Adobe Reader and Acrobat 9.3.2 and 8.2.2 to be Delivered by New Updater." Adobe: April 8, 2010. Available at http://blogs.adobe.com/adobereader/2010/04/upcoming_adobe_reader_and_acro.html.

³⁵ Betlem, Paul. "Flash Player in Chrome, An Update." Adobe: May 25, 2010. Available at http://blogs.adobe.com/flashplayer/2010/05/chrome_update.html.

The German Anti-Botnet Advice Center, announced in September 2010, is a similar public-private partnership. It feeds botnet-related network traffic data, provided by the Federal Office for Information security (BSI), to a consortium of 10 German ISPs; the ISP consortium, coordinated by a trade association named eco, the German counterpart to the IIA.⁴¹ It differs in the way it supports users, however, in that it operates a centralized notification service and support center explicitly distinct from the member ISPs themselves. Funded by the German Ministry of the Interior, the notification service and support center offers telephone-based support and provides affected users with “DE-Cleaner” — antivirus software developed by Avira, Kaspersky, and Symantec.⁴²

Japan’s Cyber Clean Center (CCC), established in 2005 as a joint venture between Japan’s Ministry of Internal Affairs and its Ministry of Economy, Trade, and Technology, detects botnet-related traffic information itself and passes the information to participating ISPs. The ISPs then notify affected users of potential infection and direct them to the CCC’s website.⁴³ The site offers users information about the process of bot infection and tools developed in-house by the CCC, in collaboration with Japan’s national CERT, to eliminate infections.⁴⁴




Figure 6. The Japan Cyber Clean Center provides common-sense computer security advice to its visitors, many of whom will have received a notice of infection from their ISPs. Here, the CCC illustrates how the use of a broadband NAT can make users less vulnerable to worms.

These initiatives usefully and positively re-frame prevailing behavioral and economic impediments to addressing botnet badware. In all three, government has assumed responsibility for disseminating available botnet-related data to market participants (ISPs) who are technically well-placed to help notify users of botnet badware infections — but who would otherwise have insufficient individual interest to carry out such notifications. Each program, moreover, reduces the overall costs associated with data sharing and badware remediation by employing government financial resources and technical expertise. By reaching out to individual users, these initiatives serve both reactive and preventative functions. Rather than allowing badware to persist on user computers indefinitely, the initiatives offer tools and resources to assist users in cleaning their computers; in doing so, they reduce the aggregate resources available to botmasters. By raising user awareness of the particularized threats badware poses, users are also more likely to engage in sound security practices.

No similar partnerships have emerged to date in the United States. Thus far, ISPs have been left to tackle botnet badware detection and remediation individually. In September 2010, Comcast expanded a trial notification program to over 16 million U.S. broadband subscribers.⁴⁵ Billing the service “Constant Guard,” Comcast notifies account subscribers of detected botnet activity by e-mail and, in some cases, injects a warning into the content of web pages as a subscriber browses the Internet; users are then directed to download a version of Norton Security Suite that Comcast provides free of charge.⁴⁶ Comcast’s decision to notify users makes it one of only two ISPs in the U.S. market to do so. In the absence of relevant incentives or a coordinated approach, other ISPs may be reluctant to invest in similar efforts.

41 Karge, Sven. “The German Anti-Botnet Initiative.” OECD: June 22, 2010. Available at <http://www.oecd.org/dataoecd/42/50/45509383.pdf>.

42 “Project Participants.” Anti-Botnet Beratungszentrum. Available at <https://www.botfrei.de/en/teilnehmer.html>.

43 “What is Cyber Clean Center?” Japan Cyber Clean Center. Available at https://www.ccc.go.jp/en_ccc/index.html.

44 “Procedure of Bot Cleaning.” Japan Cyber Clean Center. Available at <https://www.ccc.go.jp/flow/index.html>.

45 Mills, Elinor. “Comcast takes free anti-botnet service nationwide.” CNet News: September 30, 2010. Available at http://news.cnet.com/8301-27080_3-20018168-245.html.

46 “Constant Guard.” Available at <http://security.comcast.net/constantguard/>.

Takedowns and cooperation

While some stakeholders in the Internet ecosystem have responded to the badware threat by targeting botnet endpoints, others have targeted the command and control infrastructure that coordinates botnets' malicious behavior. A series of organized botnet takedown campaigns in late 2010 and early 2011 illustrates promising developments in the scope and ambition of such efforts. The campaigns' effects underscore the potential benefits of enlisting affected stakeholders through informal cooperation and legal process. Government stakeholders, moreover, showed signs of greater receptiveness to cross-border legal cooperation where cybercrime is concerned — a development that implies opportunities for broader action in the future.

Many of the challenges the security community faces today in coordinating botnet takedowns were identified in the course of the pioneering work of the Conficker Working Group. An ad-hoc collection of Microsoft employees, antivirus vendors, independent researchers, and representatives from registry operators and ICANN, the Conficker Working Group was formed in late 2008 in response to the proliferation of the Conficker worm.⁴⁷ Despite having no legal standing or formal organization, the group was able to track the evolution of the Conficker worm, capture many of the domain names attackers pointed at Conficker's command and control servers, and point the domain names to servers controlled by the group. When the criminals behind Conficker updated the worm to enable installation of arbitrary badware and use a much broader array of command and control domains, the working group coordinated with top-level domain registries and ICANN to block the registration of those domains. As a result of the group's efforts, the Conficker botnet can no longer be used for organized criminal activity. Yet despite its initial coordination, the Conficker Working Group did not engage in an organized remediation initiative: consequently, as of April 2011, computers from over 4 million IP addresses continued to exhibit signs of infection.⁴⁸

Botnet takedown efforts that have occurred in the wake of the Conficker Working Group's report have principally employed national legal processes to physically seize and sinkhole botnet command and control servers; subsequent attempts to remediate detected infections, however, have been limited by the technical and jurisdictional constraints that attend such an approach. On October 25, 2010, the Dutch High Tech Crime Team collaborated with Dutch security firm Fox-IT and the national CERT (Govcert.nl) to successfully seize 143 of the Bredolab botnet's command and control servers, which were hosted in the Netherlands. Simultaneously, the suspected botmaster was arrested with the cooperation of Armenian law enforcement.⁴⁹ While the takedown did succeed in meaningfully disrupting Bredolab's command and control operations, the scope of the operation did not encompass three additional servers hosted in Russia and Kazakhstan. The botnet, therefore, was left partially operational.⁵⁰ Furthermore, the Dutch authorities attempted to notify infected users by redirecting web browser traffic to a warning page rather than by notifying the ISPs of affected IP addresses; in doing so, the Dutch authorities probably caused confusion — and they may have violated Dutch law.⁵¹ By contrast, Microsoft's February 2011 takedown of the Rustock botnet hybridized the approaches to the Bredolab and Conficker takedowns: the takedown saw the seizure of servers corresponding to 97 IP addresses and seized thousands of domain names from U.S.-based registries.⁵² Microsoft, which is not a law enforcement body, had to establish independent standing to obtain relief through the courts; additionally, when seeking to assist ISPs in notifying customers with infected computers, it had to do so without the formal legal standing that attends requests from law enforcement, and did so in parallel with, but outside, the legal process that enabled the seizures.

⁴⁷ "Conficker Working Group: Lessons Learned." The Rendon Group: January 2011. Available at http://www.confickerworkinggroup.org/wiki/uploads/Conficker_Working_Group_Lessons_Learned_17_June_2010_final.pdf.

⁴⁸ "Conficker: World Charts." ShadowServer Foundation. Available at <http://www.shadowserver.org/wiki/pmwiki.php/Stats/Conficker#toc7>.

⁴⁹ "Dutch Team Up With Armenia for Bredolab Botnet Take Down." The New York Times: October 26, 2010. Available at <http://www.nytimes.com/external/idg/2010/10/26/26idg-dutch-team-up-with-armenia-for-bredolab-botnet-take-53590.html>.

⁵⁰ Mushtaq, Atif. "Bredolab: It's not the size of the dog in the fight..." FireEye: October 27, 2010. Available at <http://blog.fireeye.com/research/2010/10/bredolab-its-not-the-size-of-the-dog-in-fight.html>.

⁵¹ Kirk, Jeremy. "Did Dutch police break the law taking down a botnet?" Computerworld: October 26, 2010. Available at http://www.computerworld.com/s/article/9193143/Did_Dutch_police_break_the_law_taking_down_a_botnet.

⁵² Temporary Restraining Order of March 9, 2011. *Microsoft v. Does 1-11*, Case No. 2:11-cv-00222 (W.D. Wa.). Available at <http://www.noticeofpleadings.com/images/TRO.PDF>.

The most recent high-profile botnet takedown was that of Coreflood in April of 2011. The takedown effort was spearheaded by the New Haven office of the FBI: they seized command and control servers physically located in the United States, took control of domain names Coreflood used to communicate with those servers, and directed traffic from infected bots to specially designed command and control servers in the custody of the nonprofit Internet Services Consortium. The ISC then used the servers to respond to hundreds of thousands of U.S.-based infected computers with a command to shut down the malware.⁵³ The FBI has further sought judicial sanction to notify the ISPs of U.S.-based victims of the Coreflood badware; they recommended the use of a special update to Microsoft's Malicious Software Removal Tool designed to eliminate the Coreflood infection, and have requested consent from individual computer owners to issue an 'uninstall' command to the bot software itself.

Viewed as a whole, takedown efforts such as the above illustrate progressively greater degrees of ecosystem involvement in botnet takedown and remediation: they enlist security companies, national judicial systems, action by law enforcement, and notice to both ISPs and infected customers. The successes of Rustock and Coreflood are in part attributable to the fact that the United States, which is home to most commercial domain registries, can legally compel their cooperation. To date, takedowns have principally relied on botnet command and control servers' being located in a single jurisdiction; a model for cross-jurisdictional collaboration has yet to emerge.

There are signs, however, that international cooperation models are evolving to address these challenges. As an example, consider the evolving state of cybercrime enforcement actions in Taiwan. The widespread use of badware, installed in cybercafés and on home computers, to steal financial and game account credentials led to the enactment of anti-badware legislation in 2003; Taiwanese prosecutors, however, were unable to bring enforcement actions against the criminals responsible when such criminals were physically located in, or routing communications through, the People's Republic of China.⁵⁴ In April 2009, Taiwan and the PRC established a landmark agreement on criminal judicial cooperation,⁵⁵ giving Taiwanese law enforcement the ability to request assistance from the PRC in matters of badware-related cybercrime. Since then, cooperation has netted 1,329 arrests for twenty telecom and Internet fraud scams that crossed the two jurisdictions.⁵⁶

In the European Union, the European Commission has recognized the importance of establishing a framework to harmonize the laws of member states with respect to combating badware. In its September 2010 draft directive, the Commission sought to mandate that the deployment and use of badware be considered a criminal offense. The Commission also sought to liberally construe member state jurisdiction over badware-related conduct occurring within EU territory, to establish national anti-cybercrime bodies available for consultation and assistance at all times, and to set a minimum standard for response (eight hours) when urgent assistance is needed.⁵⁷

Setting standards for reporting and response

Outside of government, some participants in the security ecosystem have responded to the need for more effective badware-related data sharing by creating standards and baseline expectations for actors engaged in identifying badware and investigating badware reports. Such standards and expectations, while inherently voluntary, aim to facilitate aggregation and collection of badware-related data and, once data has been shared with relevant stakeholders, to model effective and accountable responses.

⁵³ Temporary Restraining Order of April 12, 2011. *United States v. Does 1-13*, Case No. 3:11-cv-561 (D. Conn.). Available at http://newhaven.fbi.gov/dojpressrel/pressrel11/pdf/nh041311_5.pdf.

⁵⁴ Tu, Doreen. "Cybercrimes in Taiwan: Experiences and challenges we face." Berkman Center Luncheon Series: April 5, 2011. Available at http://wilkins.law.harvard.edu/events/luncheons/2011-04-05_tu/2011-04-05_tu640.mov.

⁵⁵ "Mutual judicial assistance aims at cross-Strait harmony." Xinhua News Service: April 26, 2009. Available at http://news.xinhuanet.com/english/2009-04/26/content_11261834.htm.

⁵⁶ "Cross-strait cooperation cited as key in fight against crime: police." The China Post: March 3, 2011. Available at <http://www.chinapost.com.tw/taiwan/china-taiwan-relations/2011/03/03/293186/Cross-strait-cooperation.htm>.

⁵⁷ "Proposal for a Directive of the European Parliament and of the Council on attacks against information systems." European Commission prop. 2010/0273. Available at http://ec.europa.eu/home-affairs/policies/crime/1_EN_ACT_part1_v101.pdf.

In February 2011, after consulting with security researchers, malware remediation firms, web hosting providers, and policy advocates, StopBadware released its Best Practices for Web Hosting Providers, a high-level framework that sets expectations for web hosting providers in receipt of malware reports.⁵⁸ Web hosting providers, as we have observed, are integral parts of the Internet ecosystem; when hosting providers are notified that their services are being abused in order to deliver badware, it is crucial to set expectations for appropriate and timely mitigation and remediation actions — regardless of whether the abuse is deliberate or due to malicious compromise.

StopBadware’s best practices prescribe timely acknowledgment of badware reports; set guidelines for efficient assessment, mitigation, and remediation; define channels for informing site owners and downstream providers of the need for action; and encourage providers to examine reporting trends to improve the speed and effectiveness of remediation.⁵⁹ The best practices also provide a rubric for assessing how effectively hosting providers respond to known badware threats; providers complying with the best practices could potentially be certified as “good neighbors”, and those unable or unwilling to do so could be identified. With sufficient adoption, the practices will also help to raise webmasters’ awareness of the security threats they face and increase incentives for malware reporters to more effectively target reporting efforts.

Within the security community, efforts are underway to reduce barriers to exchanging information about badware samples for contextual and behavioral analysis. In May 2010, the IEEE Industry Connections Security Group (ICSG) released a comprehensive data format designed to facilitate such exchanges.⁶⁰ Rather than relying informally on antivirus detection signatures and narrative descriptions of a given piece of badware’s provenance, researchers can use the format to share information about a sample’s prevalence, its characteristics, and its relationship to other badware.⁶¹ Because the format is publicly available and extensible, it enables the security community to develop internal tools tailored to their organizational needs while ensuring that when data is exchanged, the tools interpret that data in a structured format.

Conclusion

Criminals promoting badware continue to take meaningful steps to increase users’ exposure to badware infection and to maximize the value of infected endpoints and distribution points to the underground economy. While any given badware attack can be evaluated based on the technical and behavioral vulnerabilities it exploits, it is clear that — at a macroscopic level — criminals are exploiting vulnerabilities in the Internet ecosystem itself. Economic models, policy and legal frameworks, approaches to user education, software development practices, and other key elements of the ecosystem are only just beginning to evolve defenses to guard against the badware threat.

We have highlighted a few examples of such evolution in action. Leading software vendors are setting a positive example for the role of applications in protecting endpoints. New initiatives are redefining the role of ISPs in keeping their customers — and the Internet more broadly — secure. Independent organizations like StopBadware are working to do the same for other sectors, such as the web hosting industry. And governments are recognizing the value of private-public partnerships and formal collaboration with other states to combat cybercrime, regardless of its source.

⁵⁸ “StopBadware releases best practices for web hosting providers.” StopBadware: March 15, 2011. Available at https://stopbadware.org/home/pr_03152011.

⁵⁹ “Best Practices for Web Hosting Providers: Responding to Malware Reports.” StopBadware. Available at <https://stopbadware.org/pdfs/Best%20Practices%20-%20Responding%20to%20Malware%20Reports.pdf>.

⁶⁰ “IEEE-SA ICSG’s Free XML Schema Brings New Efficiency to Computer-Security Industry’s Sharing of Malware Samples.” IEEE: May 24, 2010. Available at <http://eon.businesswire.com/news/eon/20100524006144/en/IEEE/industry-connections/malware>.

⁶¹ “IEEE Malware Working Group XML Schema FAQ.” Available at <http://grouper.ieee.org/groups/malware/malwg/Schema1.1/MWG-XML-Schema-faq.pdf>.

Despite these efforts, badware continues to be a substantial threat to individuals, businesses, governments, and global economy. There is a need for further work to address the many remaining weaknesses in the ecosystem. Prioritizing this work and measuring its success are both complicated by the ambiguity inherent in measuring the problem. This should be construed as an opportunity for policymakers and industry players to work together to create new methods to measure — and centralize the measurement of — key elements of the threat landscape.

Taken together, these two ideas — a shift from unilateral security to collective defense of the ecosystem, and developing a shared understanding of how to measure the problem — represent a foundation for addressing the badware threat. It is incumbent upon all of us to build tomorrow's solutions atop this foundation.