MOHEEB ABU RAJAB, LUCAS BALLARD,
PANAYIOTIS MAVROMMATIS,
NIELS PROVOS, AND XIN ZHAO

the nocebo* effect on the Web: an analysis of fake anti-virus distribution

**From the Latin, meaning "I will harm."*

Moheeb Abu Rajab is a Senior Engineer in the Infrastructure Security group at Google. His research interests are in computer and network security. He received his PhD in computer science from the Johns Hopkins University in 2008.

moheeb@google.com

Lucas Ballard is a member of the Security Team at Google. He received his PhD from The Johns Hopkins University.

lucasballard@google.com

Panayiotis Mavrommatis is a senior software engineer on Google's security team, fighting Web-based malware. His interests involve computer security and large-scale distributed systems. Panayiotis holds a Master of Engineering from Massachusetts Institute of Technology

panayiotis@google.com

Niels Provos is a Principal Engineer in the Infrastructure Security group at Google. His areas of interest include computer and network security, as well as large-scale distributed systems. He received a PhD from the University of Michigan in 2003, where he studied experimental and theoretical aspects of computer and network security at the Center of Information Technology Integration. He is the author of several popular open source libraries and security tools as well as the book *Virtual Honeypots: From Botnet Tracking to Intrustion Detection*. Dr. Provos currently serves on the USENIX Board of Directors.

niels@google.com

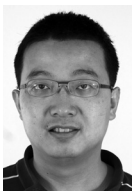

Xin Zhao, PhD, is a software engineer at Google. He holds a doctorate degree in computer science from the University of Michigan. His research interests include security, virtual machines, operating systems, and mobile systems.

xinzhao@google.com

IN RECENT YEARS, PEOPLE HAVE BECOME more aware of malware threats to their computer systems. The common advice to computer users is to install virus and malware detection. This advice has even been codified in Microsoft's Security Center, which provides prominent warnings when such protection is missing.

On the other hand, personal computer systems are lucrative targets for adversaries that compromise computers to steal and monetize sensitive information. As computer systems have become more difficult to compromise, social engineering is becoming an increasingly popular attack vector for enticing users to provide the same information without requiring any vulnerability.

Recently, a threat that we call Fake Anti-Virus has emerged. Fake AV attacks attempt to convince users that their computer systems are infected and offer a free download to scan for malware. Fake AVs pretend to scan computers and claim to find infected files—files which may not even exist or be compatible with the computer's OS. Users are forced to register the Fake AV program for a fee in order to make the fake warnings disappear. Surprisingly, many users fall victim to these attacks and pay to register the Fake AV. To add insult to injury, Fake AVs often are bundled with other malware, which remains on a victim's computer regardless of whether a payment is made.

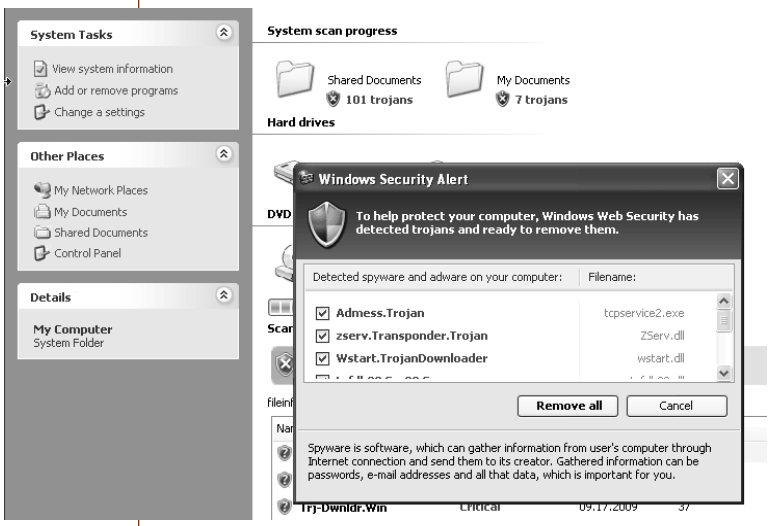

FIGURE 1: A SCREENSHOT OF A FAKE AV SITE. THE BROWSER WINDOW RESEMBLES THE LOOK AND FEEL OF WINDOWS XP.

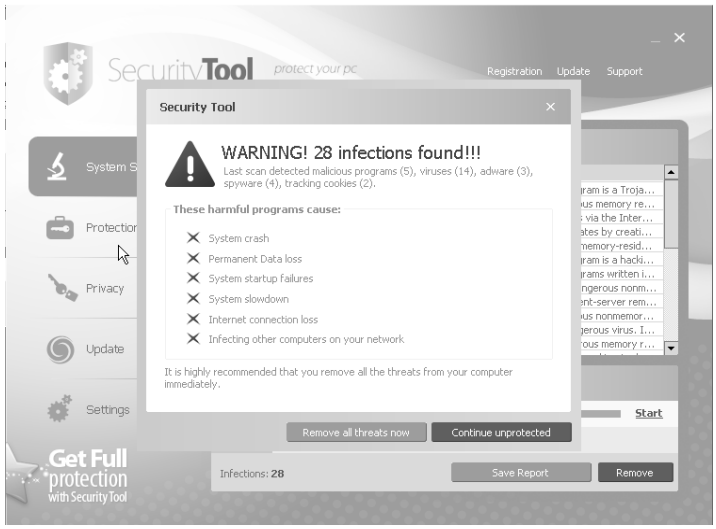

FIGURE 2: A DOWNLOADED FAKE AV BINARY WARNS OF INFECTION AND URGES THE USER TO BUY A PRODUCT.

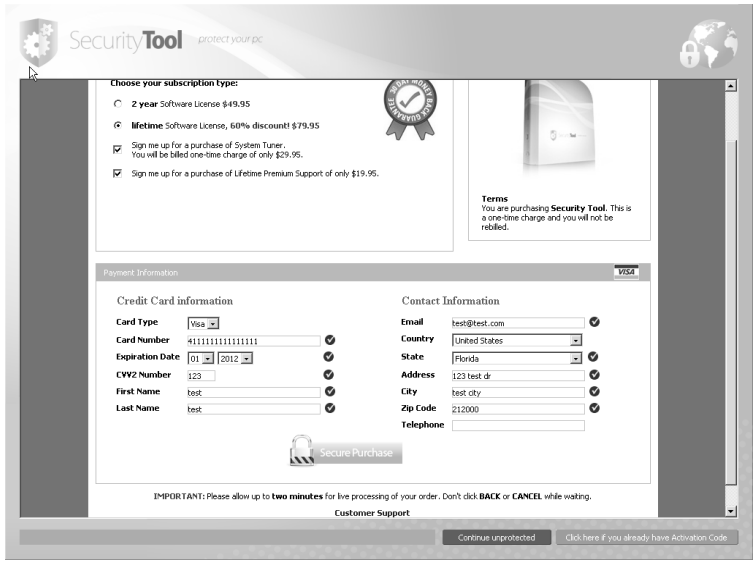

FIGURE 3: A FAKE AV PAYMENT SITE. MANY FAKE AV SITES SHARE THE SAME PAYMENT SITES.

Background

Social engineering attacks scaring users about false insecurities are not new. As early as 2003, malware authors prompted users to download Fake AV software by sending messages via a vulnerability in the Microsoft Messenger Service [5]. We observed the first form of Fake AV attack involving Web sites, e.g., malwarealarm.com, on March 3, 2007. At that time, Fake AV attacks employed simple JavaScript to display an alert that asked users to download a Fake AV executable.

More recent Fake AV sites have evolved to use complex JavaScript to mimic the look and feel of the Windows user interface. In some cases, the Fake AV detects even the operating system version running on the target machine and adjusts its interface to match. Figures 1, 2, and 3 show screenshots representative of Fake AV attacks that we frequently encounter. In Figure 1,

a Web page loads images and text that mimic the appearance of Windows XP. An animated "System scan progress" simulates an ongoing scan for viruses. This is followed by a Windows Security Alert dialog warning the user that various types of malware have been detected. At this point, the Fake AV conveniently provides the user with a button to remove the malware as shown in Figure 2. Clicking the button causes the download and installation of a Fake AV application. This application warns users that their computer is at risk, urging them to buy the full version of the software to "remove all threats." A user who chooses to purchase the software is directed to a payment site (see Figure 3) which asks for credit card information and processes the payment for registering the Fake AV software.

Discovering Fake AV Distributors

We use Google's malware detection infrastructure [2] to discover Web sites that distribute Fake AV software. Briefly, that system uses machine learning to identify potentially malicious Web pages from Google's Web repository. Each page that is flagged by the screening process is further examined by navigating to it with an unpatched Windows virtual machine running an unpatched version of Internet Explorer. Detection algorithms use signals derived from state changes on the virtual machine, network activity, and scanning results of a group of licensed antivirus engines to decide definitively whether a page is malicious.

One of the algorithms is designed to complement our licensed AV engines to specifically detect social engineering attacks, including Fake AV attacks. We do not disclose the details of the detection algorithm, due to the highly adversarial nature of this field. This algorithm is currently used to protect hundreds of millions of Web users from Fake AV attacks and disclosing it may jeopardize this effort.

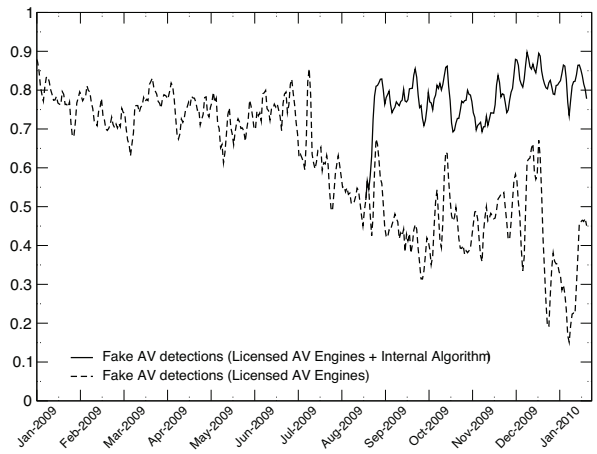

FIGURE 4: FAKE AV DETECTION RATE OVER TIME. INTERNAL ALGORITHMS COUNTER THE INCREASING ABILITY OF ATTACKERS TO EVADE AV ENGINES.

DATA COLLECTION

The data for this article was generated by reprocessing a subset of Web pages that Google's malware detection infrastructure had analyzed between January 1, 2009, and January 31, 2010. Due to the large volume of data, we only reprocessed pages that either resulted in a drive-by download, were convincingly marked as Fake AV, or were otherwise deemed "suspicious" (less than 100% confidence in a page's maliciousness) when they were first visited. Additionally, we scanned a 20% random sample of pages that were originally classified as safe. In total, we reprocessed 240 million pages to establish our data set.

We reprocessed each page using our detection algorithms and virus signatures from mid-February 2010. As Figure 4 shows, our detection rate has improved significantly, reaching up to 90% after we started using our detection algorithm in August 2009.

TERMINOLOGY

Throughout this article we use "Infection domain" to denote a domain that hosts malicious content, including exploits that cause drive-by downloads or content classified as Fake AV. Infection domains are divided into: (1) Exploit domains, which host malicious content that is not a Fake AV, and (2) Fake AV domains, which serve content that was classified as Fake AV using the aforementioned techniques.

Fake AV Distribution Is on the Rise

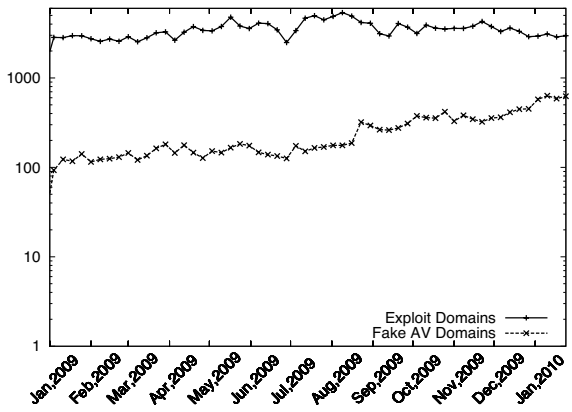

FIGURE 5: TOTAL NUMBER OF NEW INFECTION DOMAINS PER WEEK IN LOG SCALE. FAKE AV DOMAINS EXHIBIT A STEADY UPWARD TREND, WHILE EXPLOIT DOMAINS REMAIN RELATIVELY STABLE OVER TIME.

Figure 5 shows the number of unique first occurrences of both Fake AV and Exploit domains over the course of our study, aggregated by week. Clearly, there is a definitive upward trend in the number of new Fake AV domains. Exploit domains, however, remained relatively stable over time. Indeed, Fake AV accounts for an increasing share of the malware that Google discovers, rising from 3% to 15% over the course of our 13-month study.

Fake AV Domain Rotation

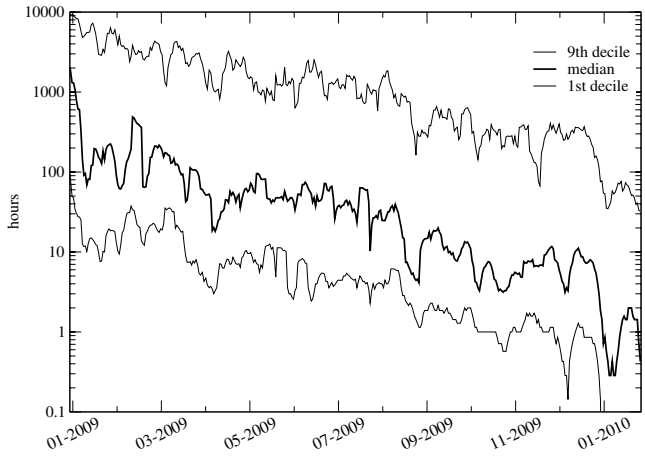

FIGURE 6: LIFETIME OF FAKE AV DOMAINS. THE MEDIAN DROPPED BELOW 10 HOURS IN SEPTEMBER 2009 AND BELOW ONE HOUR IN JANUARY 2010.

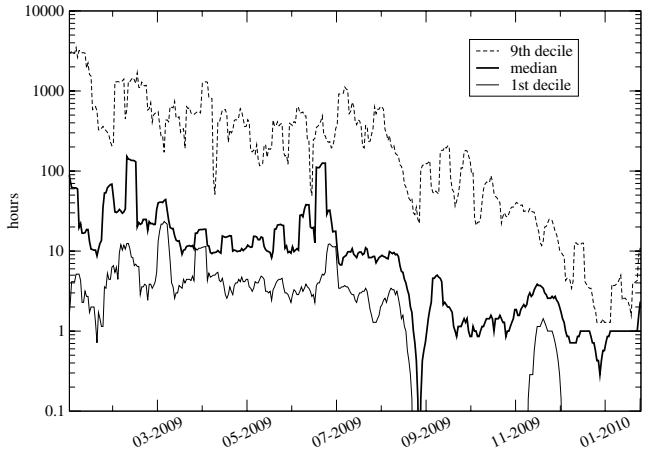

FIGURE 7: TIME TO DETECT FAKE AV DOMAINS.

Analyzing the network characteristics of Fake AV domains revealed strong affinity among groups of Fake AV domains. The 11,480 Fake AV domains mapped to 2,080 IP addresses, with 42% of these IP addresses hosting more than one Fake AV domain.

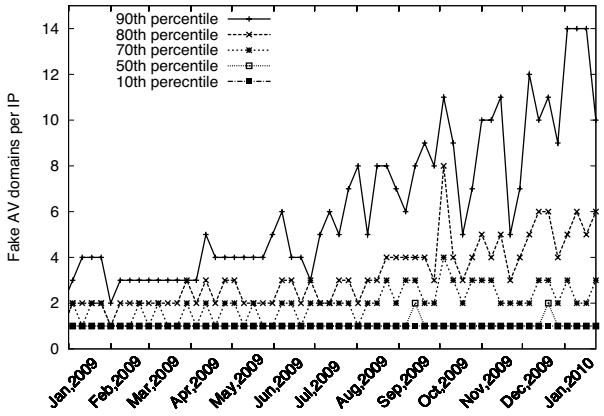

FIGURE 8: PERCENTILES OF THE NUMBER OF FAKE AV DOMAINS (OBSERVED WEEKLY) PER IP ADDRESS

Figure 8 shows an interesting trend: over time, the number of domains served from a single IP address has increased. However, as Figure 6 shows, the lifetime of these domains has actually decreased over time. These trends point to domain rotation, a technique that allows attackers to drive traffic to a fixed number of IP addresses through multiple domains. This is typically accomplished by setting up a number of domains, either as dedicated sites or by infecting legitimate sites, that redirect browsers to an intermediary site under the attacker's control. The intermediary is set up to redirect traffic to a set of active domains, which point to the Fake AV distribution servers.

Domain rotation is likely a response to domain-based detection techniques such as our Safe Browsing API [1]. In fact, we noticed a distinct correlation between our improved ability to detect Fake AVs and the observed lifetime of each domain. Figure 7 shows the trend of our detection time for these domains, measured by the interval between the time at which we would have detected the domain in our baseline data to the actual time our system added the domain to Google's Safe Browsing list. Clearly, the detection time exhibits a downward trend, reflecting an improvement in our ability to detect Fake AV domains quickly after their appearance in our data. This trend is also in line with the reduction in Fake AV lifetime, as depicted in Figure 6.

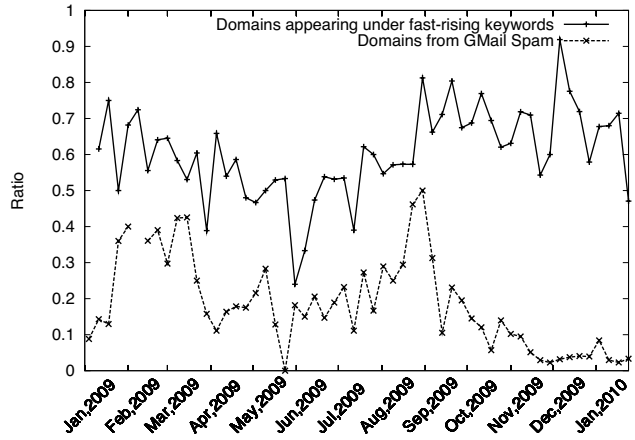

FIGURE 9: RATIO OF FAKE AV DOMAINS TO INFECTION DOMAINS AGGREGATED BY SOURCE OF THE URL. MOST INFECTION DOMAINS ENCOUNTERED ON DOMAINS THAT CONTAIN TRENDING KEYWORDS TEND TO BE FAKE AV DOMAINS.

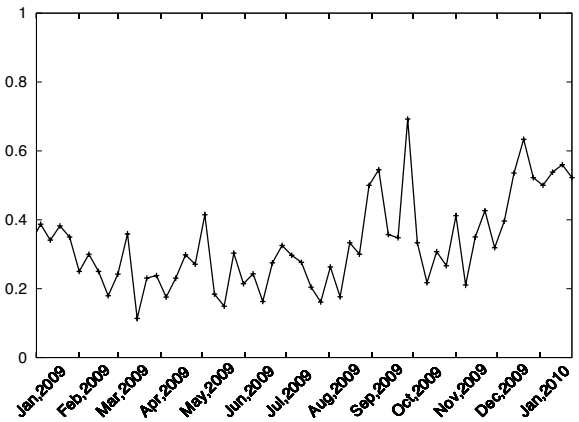

FIGURE 10: RATIO OF FAKE AV DOMAINS TO INFECTION DOMAINS ENCOUNTERED VIA AD NETWORKS. FAKE AV DOMAINS ARE EXHIBITING A RISING TREND TOWARDS AD DISTRIBUTION.

Funneling User Traffic

Fake AV distributors funnel user traffic via a set of Web sites that redirect users to the Fake AV distribution domain. We identified a number of techniques used by Fake AV distributors to lure the users into connecting to the Fake AV site: most notably, setting up dedicated spammy sites that target search engine results optimization for trendy keywords (i.e., Web-search keywords that are fast rising in popularity) and links sent directly via spam emails.

Figure 9 shows the proportion of Fake AV domains to all Infection domains when attributed to these sources. Of note, when our infrastructure identifies Infection domains on recently popular domains, 61% of the time the domain is a Fake AV domain. A smaller percentage of Fake AV domains is observed for domains first seen from Gmail spam. These results indicate that distributors of Fake AV are more successful at targeting domains associated with trending keywords than the distributors of other types of malware.

Another common infection vector for Web-based Malware is ad networks [3]. Our system encounters ad networks in two situations. First, we process URLs from Google Ads' screening pipeline to find and block malicious ads to prevent them being served to users. Second, we encounter ads from non-Google networks while processing other Web pages from Google's index. We examined our data to find Infection domains that use one or more ad networks as intermediaries. Figure 10 shows how often Fake AV domains were delivered via ad networks relative to Exploit domains. Unsurprisingly, as the popularity of Fake AV has increased, so has the number of times Fake AV domains are delivered by ad networks. What is more striking is that, even though Exploit domains are more prominent, we see approximately the same number of Fake AV domains delivered via ads as Exploit domains.

Conclusion

As users are becoming increasingly aware of the need to secure their computers, attackers have been leveraging this awareness by employing social engineering techniques to distribute Fake AV software. Our analysis of Fake AV distribution shows that Fake AV malware now accounts for 15% of all types of malware that we identify. Additionally, we find that Fake AV malware possesses interesting characteristics that distinguish it from typical

Web-based malware. For example, Fake AV domains have more Landing domains funneling user traffic than do other Infection domains.

Fake AV distributors also rely heavily on online advertisements and domains with pages that contain trending keywords. We believe that Fake AV domains have also evolved to use more agile distribution networks that continuously rotate among short-lived domains in an attempt to avoid detection. Despite continuously improving detection and mitigation techniques, Fake AV attacks persist, demanding increased awareness and broader response from the research community at large. For more information, see our publication from USENIX LEET '10 [4].

REFERENCES

[1] Google Safe Browsing API, June 2007: http://code.google.com/apis/safebrowsing/.

[2] N. Provos, P. Mavrommatis, M.A. Rajab, and F. Monrose, "All Your iFRAMES Point to Us," in *Proceedings of the 2008 USENIX Security Symposium*, pp. 1–16.

[3] N. Provos, M.A. Rajab, and P. Mavrommatis, "Cybercrime 2.0: When the Cloud Turns Dark," *Queue*, vol. 7, no. 2 (2009), pp. 46–47.

[4] M.A. Rajab, L. Ballard, P. Mavrommatis, N. Provos, and X. Zhao, "The Nocebo Effect on the Web: An Analysis of Fake Anti-Virus Distribution," in *Proceedings of the 3rd USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET' 10),* April 2010: http://www.usenix.org/events/leet10/tech/full_papers/Rajab.pdf.

[5] J. Stewart, "Windows Messenger Popup Spam on UDP Port 1026," June 2003: http://www.secureworks.com/research/threats/popup-spam/, last visited February 18, 2010.