

The Nocebo* Effect on the Web: An Analysis of Fake Anti-Virus Distribution

Moheeb Abu Rajab

moheeb@google.com

Lucas Ballard

lucasballard@google.com

Panayiotis Mavrommatis

panayiotis@google.com

Niels Provos

niels@google.com

Xin Zhao

xinzhao@google.com

Google Inc.

Abstract

We present a study of Fake Anti-Virus attacks on the web. Fake AV software masquerades as a legitimate security product with the goal of deceiving victims into paying registration fees to seemingly remove malware from their computers. Our analysis of 240 million web pages collected by Google’s malware detection infrastructure over a 13 month period discovered over 11,000 domains involved in Fake AV distribution. We show that the Fake AV threat is rising in prevalence, both absolutely, and relative to other forms of web-based malware. Fake AV currently accounts for 15% of all malware we detect on the web. Our investigation reveals several characteristics that distinguish Fake AVs from other forms of web-based malware and shows how these characteristics have changed over time. For instance, Fake AV attacks occur frequently via web sites likely to reach more users including spam web sites and on-line Ads. These attacks account for 60% of the malware discovered on domains that include trending keywords. As of this writing, Fake AV is responsible for 50% of all malware delivered via Ads, which represents a five-fold increase from just a year ago.

1 Introduction

There has been an increasing awareness of malware threats to end user computer systems. 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 such as bank log-ins and credit cards. As computer systems become more difficult to compromise, social engineering is an increasingly pop-

ular attack vector for enticing users to provide the same information without requiring any vulnerability. Phishing attacks which present content that mimics legitimate web sites have long been known as one way of stealing credentials from users. More 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.

In this paper, we use data collected from Google’s malware detection infrastructure [9] to study the prevalence of Fake AV relative to other types of web malware. Our results show that Fake AV accounts for 15% of all malware detected by our system. More troubling is the fact that Fake AV attacks spread easily without requiring any vulnerability on a victim’s computer system. Additionally, Fake AV distributors attempt to maximize their reach by posting Ads that lead to the Fake AV distribution sites, or funneling traffic through search engine-optimized web sites that are designed to rank highly for popular keywords. Our study of Fake AV distribution networks shows that Fake AV domains are becoming more agile and frequently rotate domain names. We posit that this is an attempt to combat URL based filters.

2 Background

For the following discussion, we consider a web page or binary as Fake AV if it presents content misinforming users about the security of their computers and attempts to deceive them into buying a “solution” to remove malware supposedly found during a false system

*From Latin, *nocebo*: to harm

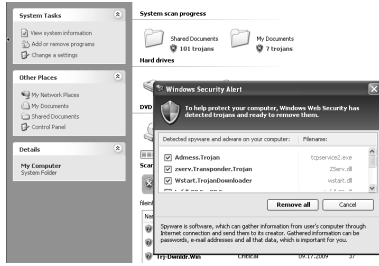


Figure 1: A screen-shot of a Fake AV site. The browser window resembles the look and feel of Windows XP.

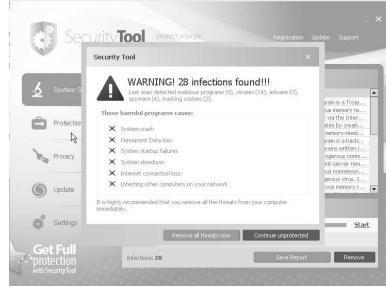


Figure 2: A downloaded Fake AV binary warns of infection and urges the user to buy a product.

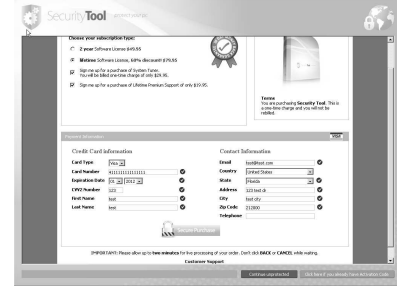


Figure 3: A Fake AV payment site. Many Fake AV sites share the same payment sites.

scan. 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 [13]. We observed the first form of Fake AV attack involving web sites, e.g. malwarealarm.com, in our systems 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 screen-shots 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 Explorer. 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. Often we find that multiple Fake AV domains use the same on-line payment service to collect registration fees.

3 Methodology

Before we present our analysis of Fake AV distribution, we describe our data set. Our analysis uses data generated by Google’s malware detection infrastructure [9]. Briefly, that system uses machine learn-

ing to identify potentially malicious web pages from Google’s repository. Each page that is flagged by the screening process is further examined by navigating to it with an un-patched Windows virtual machine running an un-patched 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 anti-virus 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 attacks¹. The algorithm is not a contribution of this work, and it suffices to mention that the algorithm is highly-resistant to false positives. As evidence, we offer that of the 22,000 malware review requests that Google received in January 2010, only 29 requests were for domains that had been flagged for distributing Fake AV. Manual inspection of these review requests revealed that our classifications were indeed correct. In this paper we classify a site as a distributor of Fake AV software if a commercial AV engine or our social-engineering detection algorithm assigns a Fake AV classification to the content downloaded from that site.

Data Collection. The data for this paper was generated by reprocessing a subset of web pages that Google’s malware detection infrastructure had analyzed between January 1, 2009, to January 31, 2010. Due to time constraints, we only reprocessed pages that either resulted in a drive-by download, were convincingly marked as Fake AV, or were otherwise deemed “suspicious”² 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.

¹We do not disclose the details of the detection algorithm due to the highly adversarial nature of the 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.

²When we visited the page there was evidence of malfeasance, but not enough to determine with 100% confidence that it was malicious.




Figure 4: Fake AV detection rate over time. Internal algorithms counter the increasing ability of attackers to evade AV engines.

We reprocessed each page using our detection algorithms and virus signatures from mid February, 2010. We argue that this method is crucial since malware-detection heuristics continually evolve, and rescanning allows us to detect web pages that may have been missed when they were originally scanned. Indeed, such trends are evident in Figure 4, which shows the original detection rates normalized against data that is reprocessed. Even when we include our internal social engineering detection algorithm, our original detection rate never exceeds 90%.

Figure 4 also shows degradation of detection rates between mid June and mid July, 2009. At this time our AV vendors were unable to detect the Fake AVs that we observed, and hence we developed our in-house detection algorithm. We believe that the reduction in detection rate was due to an increased level of polymorphism that allowed malware to evade our AV engines. Indeed, we examined trends in the number of flagged unique PE binaries that were downloaded by our VM (Figure 5, bottom), and discovered that the number of unique binaries increased from an average of 300 to 1,462 per day, causing the detection rate to plummet below 20%. While this hurt our ability to detect individual downloads at the time, we were still able to identify the unique domains that were distributing Fake AV (Figure 5, top), since most domains distributed multiple variants of their binaries. We observed other drastic dips in AV detection rate around mid August, 2009, and during the holiday season of 2009. The dip in August can be attributed to technical problems in our AV signature update pipeline, while the dip in December was likely due to lack of updates from the AV vendors. Fortunately our internal detection algorithm allowed us to weather these storms while providing protection to Google’s users.




Figure 5: Fake AV executable detection rate by AV engines. Slightly out-of-date signatures drastically reduce detection rates.

The above discrepancies further highlight the importance of using reprocessed data generated with up-to-date signatures. All other results in this paper are based on the reprocessed data. The discrepancies also provides an insight to the rate at which the perpetrators of Fake AV update their products: signatures that are only 1-2 weeks out of date can greatly reduce detection rates.

Terminology. Throughout this paper we divide domains into two groups: *Infection Domains* and *Landing Domains*. *Infection Domains* host malicious content, including exploits that cause drive-by downloads, or the HTML/JavaScript/binaries of a Fake AVs. *Landing Domains* serve web pages that causes the browser to retrieve content from *Infection Domains* without interaction from the user. Such domains could be hacked to include malicious content, or could be sites that are created with the sole purpose of distributing malware. *Infection Domains* are further divided into two groups: *Fake AV Domains* and *Exploit Domains*. *Fake AV Domains* serve content that was classified as Fake AV using the aforementioned techniques. *Exploit Domains* are all other domains that served content that exploited our VM, but did not have a Fake AV classification.

Caveats. In the following analysis we face the challenge of correctly counting domains. We decide to count sites by their domain name. That is, we use the first level host name under the top level domain. For example, if we observe Fake AV on the URL `http://host.foo.com/`, we count that as one instance of a Fake AV on the domain `foo.com/`. This means that we may conservatively group web sites that are logically distinct. For example, hosting providers that assign different users to unique hosts under the same domain. Therefore, our measurements can be viewed as a lower bound on the actual number of *Infection*

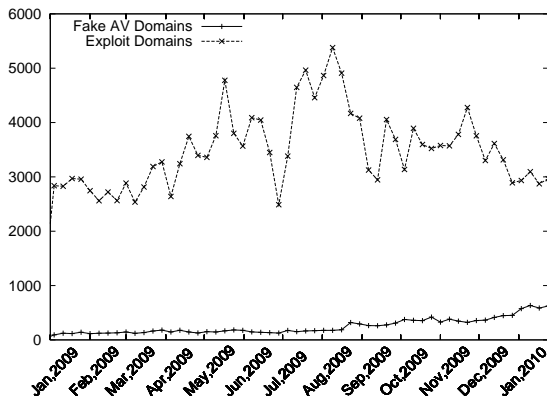


Figure 6: Total number of new Infection domains per week. Fake AV domains exhibit a steady upward trend, while Exploit domains remain relatively stable over time.

Domains. Alternatively, we could have measured Infection and Landing domains by fully qualified host names, or distinct URLs. We are concerned, however, that this could lead to over-counting Infection Domains. We reran our analysis of Landing domains (Section 4), counting at the URL granularity, and the general trends were unchanged.

4 An Empirical Analysis of Fake AVs

The goal of our work is to better understand Fake AV distribution on the Internet. In particular, we are interested in studying three high-level themes: (1) The prevalence of Fake AVs over time, both in absolute terms, and relative to other types of malware; (2) The network characteristics of domains that host Fake AV; (3) How Fake AV domains target and distribute malware.

4.1 The Rise of Fake AVs

The first goal is to measure the absolute prevalence of Fake AV domains over time. Figure 6 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 that we encounter each week. In the first week of January, 2009 we encountered only 93 unique Fake AV domains, whereas we encountered 587 in the last week of January, 2010. Interestingly, while the number of Fake AV domains increased steadily, the number of new Exploit domains fluctuated weekly, but remained relatively stable over time. Indeed, Fake AV accounts for an increasing share of the malware that Google discovers. The percent of Infection domains that were Fake AV domains increased from 3% to 15% over the course of our 13 month study.




Figure 7: Percentiles of the number of Fake AV domains (observed weekly) per IP address.

4.2 Network Characteristics

In addition to measuring the prevalence of Fake AV domains, we also wish to understand their network characteristics. In particular, we measure the locality of their hosting infrastructure, the lifetime of Fake AV domains, common relationships among different groups, and their naming conventions.

To measure the locality of Fake AV domains, we map each domain to the IP address(es) from which we successfully fetched Fake AV content. We also map each IP address to its Autonomous System (AS) using RouteViews data [11]. Overall, 11,480 Fake AV domains mapped to 2,080 IP addresses and 384 unique Autonomous Systems. This reflects strong relationships among the different Fake AV domains. We find that most of the domains were concentrated on a small number of ASes. In fact, about 52% of the ASes hosted more than one Fake AV domain, with up to 1,337 domains hosted by a single AS. Distributions across IP addresses are slightly less skewed, but even then, approximately 42% of the IP addresses hosted more than one Fake AV domain with a maximum of 334 domains detected on a single IP address. We inspected this case and found that it belongs to an ISP in Cyprus and hosted a family of domains registered under the `.info` and `.cn` TLDs. Requests were funneled to these domains via sites hosted in Russia that was spamming search engine results.

Further exploration of the number of Fake AV domains hosted per unique IP address over time (Figure 7) reveals an interesting trend: not only do multiple domains point to a single IP address, but over time, the number of domains served from a single IP address has increased. However, as the number of domains increased, the lifetime, i.e., the period over which we observed malicious content from the domain³ has actu-

³Our results are actually a lower bound on the actual lifetime of




Figure 8: Lifetime of Fake AV domains as a function of time. The median dropped below 10 hours in Sept. 2009, and below 1 hour in Jan. 2010.

ally decreased. Figure 8 shows the median, 1st and 9th deciles of the lifetime of Fake AV domains over the 13 months of our study. The lifetime decreases over time, with the median lifetime dropping below 100 hours around April, 2009, below 10 hours around September, 2009 and below 1 hour since January, 2010. These results are in-line with our findings in Figure 7.

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 Landing domains, either as dedicated sites or by infecting legitimate sites, that redirect browsers to an intermediary under the attacker’s control. The intermediary is setup to redirect traffic to a set of active domains, which point to the Fake AV distribution servers.

We hypothesize that domain rotation is a response to domain-based detection techniques. In fact, we noticed a distinct correlation between our improved ability to detect Fake AVs, and the observed lifetime of each domain. We define the time-to-detect a Fake AV domain as the interval between the time at which we would have detected the domain in our base-line data to the actual time our system added the domain to Google’s Safe Browsing list. Figure 9 shows the median, 1st and 9th deciles for the time-to-detect. Clearly, the time-to-detect 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 8.

Lastly, we measure the geographic locality of both Fake AV and Exploit domains (Table 1). These results are obtained from mapping the IP address of each Infection domain to its registrar country. Many countries had

these domains, since we do not know how long the domains served malicious content before being visited by our systems. Nonetheless, this data is useful for showing trends.




Figure 9: Time to detect a Fake AV domains as a function of time.

| Country | % Fake AV domains | Country | % Exploit domains |
|-------------|-------------------|-------------|-------------------|
| USA | 15.5 | USA | 36.4 |
| Netherlands | 14.4 | China | 13.6 |
| Germany | 11.9 | Germany | 5.0 |
| Russia | 10.9 | Russia | 3.4 |
| Canada | 9.3 | Turkey | 2.9 |
| Ukraine | 6.5 | Canada | 2.8 |
| Cyprus | 4.0 | Israel | 2.7 |
| UK | 2.8 | Netherlands | 2.7 |
| Brazil | 2.2 | UK | 2.6 |
| France | 2.2 | France | 2.3 |

Table 1: Distribution of Fake AV and Exploit domains across countries.

drastically different percentages of Fake AV domains relative to Exploit domains. For example, 13% of Exploit domains originate from China, whereas only 1.5% of Fake AV domains are located there. Additionally, Europe hosts significantly more Fake AV domains than Exploit domains.

Fake AV Domain Naming Conventions. Yet another distinguishing characteristic of Fake AV domains is the common use of deceptive naming conventions. This strategy is also widely used in phishing and scam campaigns [6], although is less common in typical Infection domains. Phishers typically embed a brand name or the name of the organization targeted by the attack in their URLs. However, unlike phishing campaigns, Fake AV domains do not seem to explicitly target any particular commercial anti-virus brand. Instead, Fake AV domains commonly use security-related English words (e.g., scan, scanner, security, anti-virus, anti-spyware, anti-malware, protect etc.). We also notice groups of Fake AV domains using very similar names with slight variations (e.g., `antimalware-softwarei0.com`, `antimalware-softwarei1.com` ..., etc.). We




Figure 10: Average number of Landing domains per Infection domain. Fake AV domains tend to infect more Landing domains than do Exploit domains.

posit that this naming convention serves two purposes: (1) it provides users with a false sense of security, and (2) it provides the Fake AV distributors with a technique to easily generate domains amenable to domain rotation.

Investigation of domains with similar names revealed that they also exhibit common network characteristics, e.g., hosted on the same IP address, used the same web server software. This indicates that Fake AV domains are affiliated with a smaller number of families that continuously change domain names. A more systematic analysis of these families is a subject of future work.

4.3 Distributing Fake AV

Finally, we explore distinguishing characteristics of Fake AV distribution by studying how Landing domains are setup to infect end users. In particular, we measure the size of Fake AV distribution networks in terms of the number of landing domains pointing to the Fake AV domain. Then we attempt to characterize how Fake AV distributors try to reach users by studying the different types of Landing domains in our data set. Finally, we study the techniques that Fake AV domains use to infect the end user.

Number of Landing Domains. We first analyze how many unique Landing domains lead to Fake AV domains, and compare it with those that lead to Exploit domains (see Figure 10). In general, a single Fake AV domain has more Landing domains (via exploiting legitimate sites, or setting up spammy sites) than a single Exploit domain, however, that ratio has been decreasing over time. In March, 2009, the ratio of Fake AV domains to Landing domains was approximately 96:1, with only 135 domains linked via 12,917 Landing domains. Conversely, the maximum ratio for Exploit domains also occurred in the beginning of March 2009, but it was only 28:1, with 2885 domains infecting 79,771 victims. In-

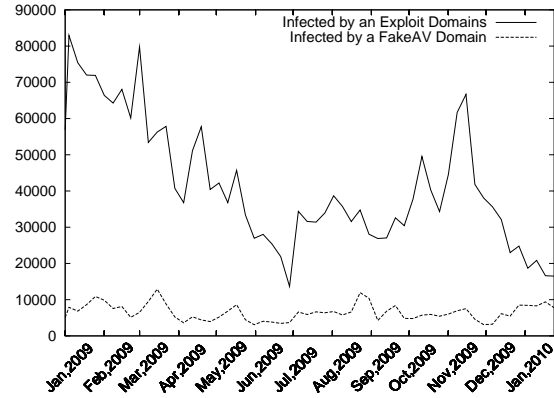


Figure 11: Total number of Landing domains classified by Infection domain.

terestingly, both Fake AV domains and Exploit domains ratios exhibit a downward trend. However, while the total number of Landing domains infected by Exploit domains has decreased significantly over time, the total number of Landing domains pointing to Fake AV domains has remained relatively stable (see Figure 11). This trend again points to domain rotation; and it seems that Fake AV domains tend to exploit this strategy more aggressively than Exploit domains in general.

Sources of Fake AV. Our infrastructure scans URLs from a number of sources, including domains that contain trending keywords (i.e., web-search keywords that are fast-rising in popularity), URLs extracted from GMail spam, and URLs from Google’s index. We examine each of these sources and aggregate Infection domains according to which source *first* included this domain. Figure 12 shows the proportion of Fake AV domains to all Infection domains when attributed to different 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.

It is important to note that the different sources are not mutually exclusive, however, they tend to be. Over the course of 13 months, we measured the overlap between the various sources. We found that for domains with fast-rising keywords, only 2.6% appeared in our GMail spam feed. Only 20% of the domains from the GMail spam feed appeared in our feed of fast-rising domains.

Another common infection vector for web-based Malware is Ad Networks [10]. Our system encounters

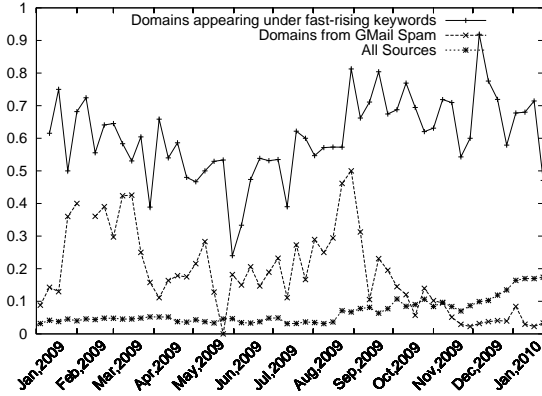


Figure 12: Ratio of Fake AV domains to Infection domains aggregated by source of the URL. Most Infection domains encountered on domains that contain trending keywords tend to be Fake AV domains.

Ad Networks in two situations. First, we process URLs from Google’s 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 13 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.

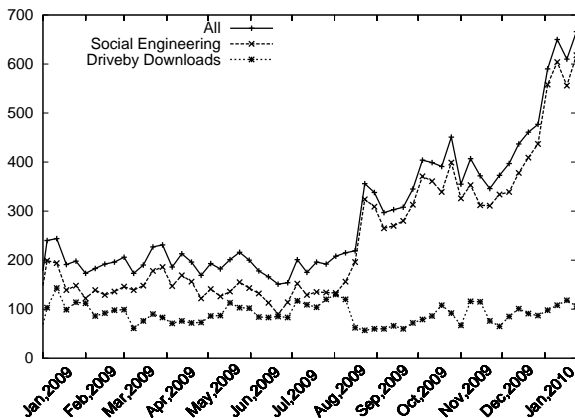


Figure 14: Number of Fake AV domains (aggregated per week) that used drive-by-download versus social engineering to deliver malware to the end user; social engineering is significantly more prevalent.

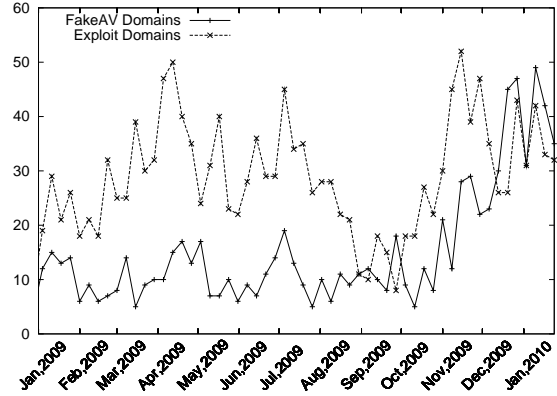


Figure 13: Total unique Infection domains encountered via ad networks. Fake AV domains are exhibiting a rising trend towards ads distribution.

Delivery Mechanisms. We analyze the mechanism that Fake AV domains use to deliver malware to the end user. In our analysis we differentiate between two cases: (1) drive-by download: in which the Fake AV malware is delivered and/or run using an exploit without requiring any user interaction, and (2) social engineering: in which user interaction was required to deliver the Fake AV. Figure 14 shows the number of Fake AV domains that used drive-by download versus those that use social engineering. Fake AV distributors predominantly use social engineering to distribute their software. While the percentage of domains that used drive-by downloads remains relatively stable, the percentage of domains that used social engineering sustained steady growth. Eventually up to 90% of all domains used social engineering. Throughout our study approximately 14% of Fake AV domains employed both drive-by downloads and social engineering.

5 Discussion

Our results are based solely on the web pages that our system observes. It is possible that some domains evaded our classification, and in such cases we may understate the threat of Fake AVs. Nonetheless, we argue that our results still provide a representative view of the recent trends and characteristics of Fake AV distribution on the web.

Due to the adversarial nature of the field we chose not to disclose the details of our detection algorithm. However, we note that our data is publicly available via Google’s Safe Browsing API [1] and the Safe Browsing Diagnostic page [2]. These services are currently used by browsers such as Google Chrome, Firefox and Safari to protect users from visiting malicious web sites. Moreover, these services can be used for comparative studies.

6 Related Work

Fake Anti-Virus is a quickly-growing attack trend. To make matters worse, Fake AV domains often target high-profile sites. For example, Facebook [12], the New York Times [5], and Twitter [8] have all been used to distribute Fake Anti-Virus (often through malicious advertisement or user posts). Indeed the severity of these attacks has even caught the eye of the United States Federal Trade Commission [4]. In December 2008 the FTC asked the United States District Court to halt Fake AV “scareware” schemes.

As the public concern for Fake AV’s has increased, so has the attention of legitimate anti-virus companies and other corporations. For example, CA published an advisory in November, 2008 describing details of a Fake AV family [3]. The advisory detailed the UI typically employed to trick users into installing Fake AV products. SecureWorks and Microsoft have also provided a general overview of Fake AVs [14, 7]. While each of these reports provide examples of potent threats—and also give potential victims tips on how to identify scams—none offers an in-depth look at the large-scale proliferation of Fake AV.

While there has been no peer-reviewed literature on the general trends of Fake AV, such behavior was first mentioned by Provos et al. [10]. They reported finding 60 Fake AV domains per week from July to October 2008, and 148,000 URLs funneling traffic to 450 Fake AV domains in January 2009. This paper expands upon those initial observations and goes into greater depth on the trends of Fake AV distribution. Symantec also studied several Fake AV campaigns over the course of two months in late 2009 [15]. Their study uses samples captured by their AV scanners to examine the financial motivation of this type of malware, and to provide an in-depth analysis of several different examples of Fake AV. They also briefly studied some Fake AV domains (although it was unclear from the discussion how they measured a domain).

Instead of focusing on the financial aspects and specific examples of Fake AV, this paper provides a 13 month study of Fake AV distribution on the Internet. We examine how Fake AV domains differ from typical Exploit domains, focusing on the techniques used to reach a large number of users while attempting to evade detection. Additionally, we leverage Google’s vantage point and use a reprocessed data-set to establish a level of ground-truth for our analysis.

7 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. Unfortunately, many users fall victim to these attacks and are tricked into paying a phony subscription fee and into divulging personal information. In this paper, we analyze Fake AV attacks by studying 13 months of reprocessed data from Google’s malware detection infrastructure. This data 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 distinguishes 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 on-line 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 continue to persist, demanding increased awareness and broader response from the research community at large.

Acknowledgements

We would like to thank Oliver Fisher, Nav Jagpal, Fabrice Jaubert, Pierre Phaneuf and Ke Wang for their contributions to Google’s malware detection infrastructure. We also thank our anonymous reviewers for their insightful comments.

References

- [1] Safe Browsing API, June 2007. See <http://code.google.com/apis/safebrowsing/>.
- [2] Safe Browsing diagnostic page, May 2008. See <http://www.google.com/safebrowsing/diagnostic?site=yoursite.com>.
- [3] Virus Detail: Win32/FakeAV Family, November 2008. <http://www.ca.com/securityadvisor/virusinfo/virus.aspx?id=74100>, last visited February 13, 2010.
- [4] Federal Trade Commission. Court halts bogus computer scans, December 2008. <http://www.ftc.gov/opa/2008/12/winsoftware.shtm>, last visited February 13, 2010.
- [5] R. Ferguson. New York Times pushes Fake AV malvertisement, September 2009. <http://countermeasures.trendmicro.eu/new-york-times-pushes-fake-av-malvertisement/>, last visited February 13, 2010.
- [6] S. Garera, N. Provos, M. Chew, and A. D. Rubin. A framework for detection and measurement of phishing attacks. In *WORM '07: Proceedings of the 2007 ACM workshop on Recurring malware*, pages 1–8, New York, NY, USA, 2007. ACM.
- [7] Microsoft Security. Watch out for fake virus alerts. <http://www.microsoft.com/security/antivirus/rogue.aspx>, last visited February 13, 2010.

- [8] B. Prince. Rogue Twitter Accounts Blasting Out Links for Fake Antivirus, September 2009. http://securitywatch.eweek.com/twitter/rogue_twitter_accounts_blasting_out_dangerous_urls.html, last visited February 13, 2010.
- [9] N. Provos, P. Mavrommatis, M. A. Rajab, and F. Monroe. All your iframes point to us. In *USENIX Security Symposium*, pages 1–16, 2008.
- [10] N. Provos, M. A. Rajab, and P. Mavrommatis. Cybercrime 2.0: When the cloud turns dark. *Queue*, 7(2):46–47, 2009.
- [11] David Meyer, University of Oregon RouteViews Project. <http://www.routeviews.org/>.
- [12] C.-A. Skinner. Fake Antivirus Scam Hits Facebook, January 2010. http://www.pcworld.com/article/188147/fake_antivirus_scam_hits_facebook.html, last visited February 13, 2010.
- [13] 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.
- [14] J. Stewart. Rogue antivirus dissected, October 2008. <http://www.secureworks.com/research/threats/rogue-antivirus-part-1/?threat=rogue-antivirus-part-1>, last visited February 13, 2010.
- [15] Symantec, Inc. Symantec report on rogue security software. White paper, Symantec Enterprise Security, October 2009. http://www4.symantec.com/Vrt/wl?tu_id=XuOB125692283892572210, last visited February 13, 2010.