



FireEye Advanced Threat Report – 1H 2011

How advanced attacks succeed, despite \$20B spent annually on enterprise IT security




Table of Contents

Inside This Report _____	1
Executive Summary _____	1
Finding 1: 99% of enterprise networks have a security gap despite \$20B spent annually on IT security. _____	2
Finding 2: Successful attacks employ dynamic, “zero-day” malware tactics. 90% of malicious binaries and domains change in just a few hours; 94% within a day. _____	4
Finding 3: The fastest growing malware categories are Fake-AV programs and Info-stealer executables. _____	6
Finding 4: The “Top 50” of thousands of malware families generate 80% of successful malware infections. _____	8
Conclusions _____	10
Methodology _____	11

Inside This Report

This is not your typical threat report. You won't see a tally of the millions of well-known malware or billions of spam messages. Those statistics appear in standard security reports that deal primarily with well-known, well-understood threats. The standard security reports have focused mainly on the threats that have been around for months or years and have been published by traditional security vendors of firewall, IPS, antivirus, and Web/email gateways.

To complete the threat landscape picture, the FireEye Advanced Threat Report focuses on the threats that have successfully evaded traditional defenses. These are the unknown threats and advanced attacks that are dynamic, targeted, and stealthy. And, they are extremely effective for compromising organizations' networks. FireEye is in the unique position to illuminate this advanced cyber attack activity since our appliances are deployed in enterprises across the globe as the last line of network defense behind firewalls, IPS, and other security gateways. Given this unique position, we are able to share our findings on the advanced threats that routinely bypass signature-, reputation- and basic behavior-based technologies, the \$20B spent on IT defenses each year.

This report dives into the FireEye Malware Intelligence Labs' analysis of shared threat data from global deployments of FireEye Malware Protection Systems (MPS). This threat data is anonymized, real-time information shared by brand-name enterprises, government agencies and educational institutions that subscribe to our Malware Protection Cloud (MPC).

Executive Summary

Based on our analysis of 1H2011 threat data, today's cyber criminals are breaking through traditional security defenses at an alarming rate despite the \$20B invested in IT security in 2010. We are clearly in a new era of dynamic cyber attacks that are very successful at evading traditional defenses, leaving virtually every enterprise vulnerable to data theft, cyber-espionage and intellectual property alteration, theft and destruction.

Traditional defenses continue to rely too heavily on signatures, reputation and lightweight behavior heuristics as well as clinging to an outmoded viewpoint on threat detection; namely focusing on known attacks while ignoring unknown attacks and associated callback channels. Based on the threat data we reviewed, criminals have developed workarounds to bypass traditional defenses using dynamic code as well as utilizing sophisticated social engineering to fool even the most educated users.

As the findings below show, to close the gap in their networks, enterprise security leaders must assume that their networks are compromised, familiarize themselves with the nature and intent of modern attacks, and supplement the traditional defenses they currently use with tools designed for today's sophisticated attacks.

Key Findings

- 1) 99% of enterprises have a security gap, despite \$20B spent annually on IT security. Within a given week, the typical enterprise network has anywhere from hundreds to thousands of new malicious infections and all industries are under sustained attack.
- 2) 90% of malicious executables and malicious domains changed in just a few hours. The dynamic nature of modern attacks is the primary means to bypass signature-based tools, making defenses such as antivirus and URL blacklists ineffective.
- 3) The fastest growing malware categories are Fake-AV programs, which take part in extortion tactic and info stealers, which abscond information.
- 4) The top 50 out of thousands of malware families account for 80% of successful infections. Sophisticated toolkits and other means are enabling the rapid production of advanced malware.

Finding 1: 99% of enterprise networks have a security gap despite \$20B spent annually on IT security.

Despite the massive investment in IT security equipment each year, our analysis of FireEye MPS deployments shows that essentially all enterprises are compromised with malware: 99% of enterprises had malicious infections entering the network each week, and 80% of enterprises faced more than one hundred infections per week, with many in the thousands per week. The median weekly infection caseload was 450 infections per week (normalized per Gbps of traffic)—with wide variations (Figure 1).

These are all events that have made it through standard gateway defenses, such as firewalls, next-generation firewalls, IPS, antivirus, email and web security gateways. These malicious events make it




Figure 1: Widespread Infections Seen Across the Enterprise

through because traditional security systems either rely on signatures, reputation and crude heuristics or were originally designed for policy control. They no longer keep up with the highly dynamic, multi-stage attacks that have become common today for targeted and APT attacks.

Note: FireEye detects malware already active within the network, as well as new malware attacking the network. FireEye systems can be deployed inline with active blocking to act as an effective counter-measure for these types of incidents.

Even the most security-conscious industries are fraught with dangerous infections.

Every company studied in every industry looks to be vulnerable and under attack. Even the most security-conscious industries, such as Financial services, health care and government sectors, which have intellectual property, personally identifiable information, and compliance requirements—show a significant infection rate.

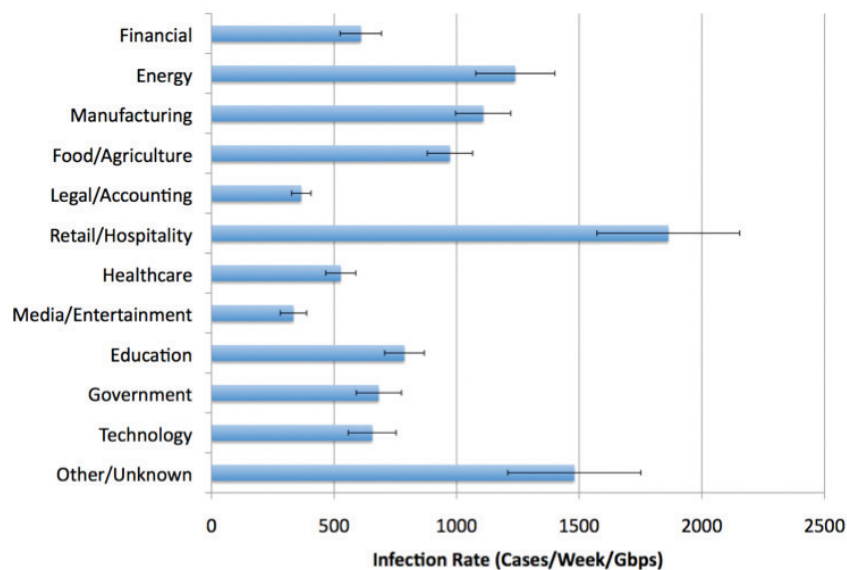


Figure 2: Average infection rates of enterprises by major vertical market segment (cases per week per Gbps of aggregate egress traffic). Note that the long-tailed nature of the distribution of infection rates could make some of these estimates noisy.

Based on this data, we see that today's cyber criminals are nearly 100% effective at breaking through traditional security defenses in every organization and industry, from security-savvy to security laggards.

Today's attacks also exhibit a global footprint with infected sites, malicious servers, and callback destinations distributed around the world.

FireEye's Global Malware Tracker




Figure 3: Global look at the infection landscape with callbacks in red indicating the data exfiltration destination, green represents CnC servers, and blue towers are bots.

Finding 2: Successful attacks employ dynamic, “zero-day” malware tactics. 90% of malicious binaries and domains change in just a few hours; 94% within a day.

Our Q2 2011 data showed that 90% of both malicious binaries (MD5 hash files) and malicious domains (URLs hosting malware) changed almost immediately, and 94% changed within a day. This dynamism increased noticeably from Q1 to Q2 2011.

We believe the daily morphing of malicious binaries and domains is timed to stay ahead of the typical practice of daily DAT and blacklist/reputation updates, enabling the malware to remain undetected and its communications unblocked.

Those that change within a few hours stay ahead of centralized “real-time” threat intelligence services that assess risk based on signatures, reputation, and behavior. Those that change once a day stay ahead of defenses that use scheduled daily updates.

We have visibility into these tactics because, unlike traditional defenses, FireEye uses a unique virtual execution (VX) engine to inspect and confirm malware, isolate callbacks and take appropriate action. We fully execute the code to see what it does, rather than assessing risk based on assumptions, signatures or formulae.

Figure 4 shows the distribution of observed lifetimes of malicious binaries, the period from the first time we saw a given MD5 signature for a particular executable to the last time we saw it across our sample.




Figure 4: Cumulative distribution function for the period over which we observe a particular binary that is the exact sequence of bytes as captured by the MD5 signature.

Malicious executables are constantly being repacked to appear new each time. Most of the MD5s we observe are so dynamic that they persist for an hour or less or are seen just once. The curve has moved noticeably up and to the left from Q1 to Q2, indicating that a smaller fraction of malware samples remain unchanged over the course of days (note that this is despite the fact that the Q2 sample is larger than the Q1 sample, increasing the size of our view into malware behavior). It's also striking that the curve steps up at each 24-hour interval indicating that some malware authors are using an integer number of days as the expiration time before they generate a new packing.

Next, we look at the malicious URLs that led to a compromise of the operating system, and extract the domain name in that URL. In Figure 5, the data shows the length of time the domains stayed unchanged.




Figure 5: Cumulative distribution function of the observed malicious domain lifetime from Q1/2011 and Q2/2011 data

Malicious domain names are changing on sub-day timescales, and the situation has worsened since the beginning of 2011.

Note that we are not implying that all malware attacks are dynamic, just that the successful attacks penetrating through the signature and reputation-based defenses use dynamic tactics to defeat those static defenses. Finding 1 showed that such threats exist on an important scale, and Finding 2 demonstrates that these threats are rapidly getting more dynamic.

Therefore, we believe that dynamic binaries and dynamic domains form the core of today's advanced, zero-day malware tactics. Cybercriminals are moving quickly and building maneuverability into their tools and operations.

In part, the move to malware dynamism explains the rapid expansion in botnets. For example, criminals need more IP addresses (aka bots or zombies) to evade signature- and reputation-based filters.

Another conclusion from these findings is that network defenses must tool up for constant change and resilience. Countermeasures must be designed for highly dynamic threats across vectors, such as Web and email. We also see a trend in which organizations must treat every attachment or Web object as suspicious.

Finding 3: The fastest growing malware categories are Fake-AV programs and Info-stealer executables.

While malware programs have multiple capabilities, the FireEye research team provides a general categorization of each malware executable with what we believe to be its primary purpose. For example, Click Fraud software makes money by creating automated HTTP transactions to particular websites in the interest of distorting (driving up) payments to advertisers. Fake-AV software is sold on the pretense that it has found non-existent malware on consumer computers and then offering to "clean" out the infection if consumers buy the full version.

Figure 6 shows the distribution of these primary malware purposes across our sample in the first two quarters of 2011.




Figure 6: Top categories of malware in the first two quarters of 2011 as labeled by FireEye Labs.

Several things stand out. The three largest categories of malware in Q2 are Fake-AV (listed as Rogue Anti_malware), Downloader Trojans (whose primary function is to download other pieces of malware), and information stealers of various forms. Comparing to Q1, we see a striking growth in Fake-AV (Rogue Anti_malware) and information stealing malware most likely due to a successful monetization model.

Of these, the information stealers are clearly the greater threat to corporate integrity. While we would certainly not advocate ignoring Fake-AV programs—they are a threat to employees' private finances and act as a conduit for more serious malware infections—it's clear that information theft is currently the highest priority problem for enterprises.

Within the category of Information Stealers, Figure 7 shows the most widespread types of malware that we observe. Readers are cautioned that, particularly for information theft, the most widespread and the most serious may be very different.




Figure 7: Top information stealers observed in our sample in Q2 2011.

- **Zbot (Zeus)** Primarily a banking Trojan, Zbot has become extremely famous for fraud against online banking for both consumers and small and medium enterprises and likely represents a high priority threat even to large enterprises in the form of fraud against senior executives.
- **Papras (aka Snifula)** has received far less publicity, but in our sample it appears to have become just as widespread as Zbot. Papras is less specialized: it steals account credentials for various online services and also logs information entered in web forms. As such, it's probably a basic tool in a number of different kinds of manually directed intrusions and information thefts.
- **Zegost** is also primarily a keylogger
- **Multibanker** are specialized banking trojans.
- **Coreflood** is a botnet that operated in many versions for ten years until taken down by the Department of Justice in April of 2011
- **Licat** is believed to be associated with Zbot.

Finding 4: The “Top 50” of thousands of malware families generate 80% of successful malware infections.

In reviewing several hundreds of thousands of events, we found that the vast majority of them derive from a few hundred malware families (as evidenced by the particular callback protocol we detected in use), and that the Top 50 most frequent malware families are represented in about 80% of all cases. Figure 8 illustrated the fraction (y) of all our cases explained by the top x malware callbacks.




Figure 8: Cumulative fraction of all cases accounted for by malware callback protocol through a particular rank in popularity. The top 20 callbacks showed up in about half of cases, and the top 300 accounted for 98+%. There is a long tail out into thousands of more rare malware families. The shape of this curve changed only slightly from Q1 to Q2.

From the figure, we can see that the exploding zoo of malware executables can be attributed to a much smaller number of malware toolkit code bases. In reviewing the top 50 families, the more successful code bases have optimized aspects of their malware binary output to be dynamic and deceptive.

Note that the frequency of appearance is not correlated with risk. One of the most common malware families, Fake-AV, extorts payments from users for falsified virus scans. This class of malware is less of a concern from an enterprise perspective, though Fake-AV should be seen as a “gateway malware” to introduce more serious information-theft malware into the network. On the other hand, nation-state APT malware used for espionage is likely to be out in the long tail of comparatively rare malware. In the range between these two zones, we find very potent, very dangerous attacks.

Many of the Top 50 attacks reflect advanced malware used by criminal syndicates for financial gain. This variety of threat is characterized by periodic campaigns combining exploit toolkits and specific malware families such as “Rogue AV” or “Fake-AV.” The attacks cast a relatively “wide but shallow” net, harvesting data and relying on automation for efficiency and profitable success rates.

Here's the anatomy of a typical "wide and shallow" attack, one that is dynamic and short-lived (in each campaign), but not especially targeted or heavily personalized:

- Hunt new victims for a few hours at certain infectious IP addresses
- Install malware via drive-by download or phishing campaign (possibly run through a social networking site)
- Collect account data from victims' computers (or install data-stealing malware on these hosts)
- Pause (or move on to a new site)
- Monetize the data that has been collected (for perhaps days or weeks)
- Run another campaign with a tweaked version of the malware and different IP addresses when we look at malware by family, and the event timeline of malware activity, we see evidence of the compressed timelines used in campaigns today. We see sharp spikes. Even with a relatively protracted activity, like that shown with Rogue.AV, we see significant spikes above a significant baseline. Note that Bot.Conficker.B remains quite active. The following Fake AV families were in the top 10: Rogue.AV, Rogue.FakeRean and Rogue.FakeAV.




Figure 9: Fraction of all cases involving the top 20 malware callbacks detected by day within 1H of 2011. Note the extremely spiky nature of the data suggests a series of short campaigns by particular groups rather than ongoing steady activity.

The other major category of attack is the "Narrow and Deep" attack that includes targeted and APT attacks. These attacks infect a relatively small number of machines that act as the beachhead from which to further infiltrate other enterprise systems, especially those that contain critical or sensitive information. The deeper infiltration is accomplished via lateral movement by propagating the malware infection to other systems and servers in the enterprise network. Only real-time monitoring of suspicious code will detect these subtle attacks.

How do criminals make their malware and domains dynamic? Point-and-click Toolkits

Criminals make code appear new by packing, encrypting, or otherwise obfuscating the nature of the code. Malware toolkits like Zeus (banking Trojan) and Blackhole (drive-by downloads) automate this process today, which we believe explains some of our finding of increasing and almost ubiquitous dynamism.

The prevalence of dynamic domain addresses indicates that criminals are moving their distribution sources very quickly as well, like a drug dealer moving to a different street corner after every few deals. By moving their malware to an unknown site (often a compromised server or zombie), and using short URLs, cross-site scripting or redirects to send traffic to that site, the criminals can stay ahead of reputation-based defenders.

Criminals invest in toolkits and dynamic domains because signatures and reputation engines have become adept at blacklisting known bad content and “bad” or “risky” URLs sites. Any stationary criminal assets will quickly be blacklisted, therefore these assets must move to remain valuable.

Conclusions

The new breed of cyber-attacks are evading existing defenses by using dynamic malware, toolkits and novel callback techniques, leaving virtually every enterprise vulnerable to data theft and disruption. Although enterprises are investing \$20B per year on IT security systems, cybercriminals are able to evade traditional defenses, such as firewalls, IPS, antivirus and gateways, as they are all based on older technology: signatures, reputation and crude heuristics.

Enterprises must reinforce traditional defenses with a new layer of security that detects and blocks these sophisticated, single-use attacks. New technologies are needed that can recognize advanced malware entering through Web and email, and thwart attempts by malware to call back to command and control centers. This extra defense is designed specifically to fight the unknown threats, such as zero-day and targeted APT attacks, thereby closing the IT security gap that exists in all enterprises.

Methodology

The analysis in this report is based on observations by FireEye Web Malware Protection System deployments, which detects inbound Web attacks and multi-protocol malware callbacks. The data in this report was obtained from customers during the period of 1H 2011. The sample size was in the hundreds and were drawn from mainly large and medium-sized enterprises and from many different vertical segments.

Frequently we may see many symptoms of malware on a given infected client: the inbound exploitation, multiple malicious binaries being downloaded, and then callback evidence of multiple malware families. Often, to become infected with one piece of malware is to become infected with many pieces. For the purpose of this analysis, we aggregate all evidence of malware that we have on a given client IP address into an "infection." That infection is the unit of analysis throughout this report. If a given IP address shows no symptoms for seven consecutive days, we consider that infection closed and any further symptoms will count as a new infection.

All the usual caveats apply here: we are observing complex enterprise networks, of unknown topology, typically from the egress points where such networks touch the Internet. Our infection counts could be off due to DHCP lease expirations that do not preserve IP address on release, physical moves of equipment, particularly laptops, presence of multiple systems behind internal NAT devices, etc.

About FireEye, Inc.

FireEye is the leading provider of next-generation threat protection focused on combating advanced malware, zero-day and targeted APT attacks. FireEye's solutions supplement security defenses such as next generation and traditional Firewalls, IPS, AV and Web gateways, which can't stop advanced malware. These technologies leave significant security holes in the majority of corporate networks. FireEye's Malware Protection Systems feature both inbound and outbound protection and a signature-less analysis engine that utilizes the most sophisticated virtual execution engine in the world to stop advanced threats that attack over Web and email. Our customers include enterprises and mid-sized companies across every industry as well as Federal agencies. Based in Milpitas, California, FireEye is backed by premier financial partners.