

May 29th, 2:40 PM

Investigative Techniques of N-Way Vendor Agreement and Network Analysis Demonstrated with Fake Antivirus

Gary Warner

The University of Alabama at Birmingham, gar@uab.edu

Mike Nagy

The University of Alabama at Birmingham, mikenagy@uab.edu


Kyle Jones

The University of Alabama at Birmingham, kjones23@uab.edu

Kevin Mitchem

The University of Alabama at Birmingham, kmitchem@uab.edu

Follow this and additional works at: <http://commons.erau.edu/adfsl>

 Part of the [Aviation Safety and Security Commons](#), [Computer Law Commons](#), [Defense and Security Studies Commons](#), [Forensic Science and Technology Commons](#), [Information Security Commons](#), [National Security Law Commons](#), [OS and Networks Commons](#), [Other Computer Sciences Commons](#), and the [Social Control, Law, Crime, and Deviance Commons](#)

Scholarly Commons Citation

Warner, Gary; Nagy, Mike; Jones, Kyle; and Mitchem, Kevin, "Investigative Techniques of N-Way Vendor Agreement and Network Analysis Demonstrated with Fake Antivirus" (2014). *Annual ADFSL Conference on Digital Forensics, Security and Law*. 3.
<http://commons.erau.edu/adfsl/2014/thursday/3>

This Peer Reviewed Paper is brought to you for free and open access by the Conferences at Scholarly Commons. It has been accepted for inclusion in Annual ADFSL Conference on Digital Forensics, Security and Law by an authorized administrator of Scholarly Commons. For more information, please contact commons@erau.edu.

EMBRY-RIDDLE
Aeronautical University,[™]
SCHOLARLY COMMONS

(c)ADFSL



INVESTIGATIVE TECHNIQUES OF N-WAY VENDOR AGREEMENT AND NETWORK ANALYSIS DEMONSTRATED WITH FAKE ANTIVIRUS

Gary Warner
gar@uab.edu

Mike Nagy
mikenagy@uab.edu

Kyle Jones
kjones23@uab.edu

Kevin Mitchem
kmitchem@uab.edu

The University of Alabama at Birmingham
Birmingham, AL

ABSTRACT

Fake AntiVirus (FakeAV) malware experienced a resurgence in the fall of 2013 after falling out of favor after several high profile arrests. FakeAV presents two unique challenges to investigators. First, because each criminal organization running a FakeAV affiliate system regularly alters the appearance of their system, it is sometimes difficult to know whether an incoming criminal complaint or malware sample is related to one ring or the other. Secondly, because FakeAV is delivered in a “Pay Per Install” affiliate model, in addition to the ring-leaders of each major ring, there are many high-volume malware infection rings who are all using the same malware. Indeed, a single criminal could participate in multiple affiliate programs using the same spreading and distribution system. Because of this, traditional malware clustering may identify common code, but fail to achieve distinction or attribution of the individual affiliate actors profiting from the scam. By combining n -way vendor agreement and live network capture, malware samples can quickly be associated with particular affiliate infrastructure and/or managing affiliate programs, while identifying and helping to prioritize investigations.

1. INTRODUCTION

Fake Antivirus is a form of Crimeware. The Anti-Phishing Working Group defines Crimeware as “software that performs illegal actions unanticipated by a user running the software, which are intended to yield financial benefits to the distributor of the software” (APWG, 2006). FakeAV malware is sometimes called “Scareware” in that the method of earning revenue is to convince the victim that their computer is infected with malware and that their only hope of removal is to buy the advertised product (Samosseiko, 2009) (Federal Trade Commission v. Innovative Marketing, Inc., 2008). The “Fake” in FakeAV comes from the fact that the purchased product provides no protection at all. In the world of FakeAV, the malware that an individual becomes infected with is installed by an affiliate. In FakeAV malware affiliate programs, centralized cyber criminals go to the expense and effort to design a sleek user interface, provide program functionality, and manage the billing and invoicing. The job of the affiliate is to get the malware to execute on as many user computers as possible. Some do this via spam, and others through “drive-by” installers, which covertly execute when someone visits a website that has been compromised (John, Yu, Xie, & Abadi, 2011) or views a malicious advertisement (Provos, Mavrommatis, Rajab, & Morose, 2008). Some of these affiliate

programs in the past have included Gagarincash (2011), Gizmo, Nailcash, Best AV, Blacksoftware, and Sevantivir.com (Krebs, 2011). University of California, Santa Barbara (UCSB) studied records from three FakeAV affiliate programs and documented 106 million successful infections which led to 2.2 million purchases of the FakeAV software generating \$133 Million dollars in revenue for the criminals (Stone-Gross, et al., 2013).

2. A HISTORY OF FAKE AV CASES

In 2010, three arrests were made by the Federal Bureau of Investigation (FBI) in a Fake Antivirus case known as Innovative Marketing. Millions of computers were infected with Scareware after viewing malicious advertisements placed by the Innovative Marketing crew under many fake names (Warner, 2008). The Federal Trade Commission (FTC) civil case against Innovative Marketing, which began in 2008, concluded in June 2013 with a \$163 Million judgment against Kristy Ross, the leader of Innovative Marketing (FTC, 2013). Ross' ads were displayed more than 600 million times and more than one million victims purchased the fraudulent software.

In June of 2011, the FBI worked with seven national law enforcement agencies around the world in a coordinated cybercrime effort, resulting in arrests and seizure of computers in France, Germany, Latvia, Lithuania, the Netherlands, Sweden, Ukraine, the United Kingdom, and the United States. The gang of cybercriminals was accused of having stolen \$72 million by tricking more than 960,000 victims into buying Fake Anti-virus products with their Scareware technique (FBI Press, 2011).

Five years after the FTC case against Ross began, FakeAV is still being used to infect home computers via malicious advertisements. As recently as January 7, 2014, the "Daily Motion" website was hosting advertisements that would cause visitors to have Scareware installed on their computer (Mimoso, 2014).

In Microsoft's most recent Security Intelligence Report (SIR), covering the first half of 2013, Fake AV is described as "Rogue Security Software" and that it "has become one of the most common methods that attackers use to swindle money from victims" (Microsoft, 2013). The same report gives Microsoft's names for the most prominent Rogue Security Software seen in the previous 12 months, making it clear that Win32/Winwebsec and Win32/FakeRean were by far the most dominant versions recently seen.

3. CONFUSION OF MALWARE NAMES & CASES

From a law enforcement perspective, the distinction between Winwebsec and FakeRean is only useful if it were true that all of the malware belonging to the Winwebsec group is being operated by one criminal organization and all of the malware belonging to the FakeRean group is being operated by another criminal organization. While there are some types of malware botnets where it is true that the entire botnet is operated by one commercial entity, it is often true that there are many competing criminals participating in the same space, often using common, shared, or stolen source code as a starting point for new variants of malware. At one extreme are the single-controller botnets, such as the Peer to Peer (P2P) Torpig botnet explored by UCSB (Brett Stone-Gross, 2009), or recent versions of P2P Zeus botnets analyzed by CERT Polska (CERT Polska, 2013). On the other end are malware packages that are intended to be sold as stand-alone infection and management kits, most famously Poison Ivy analyzed by Paul Rascagnères of Malware Luxembourg (Rascagnères, 2013), where every hacker buys a private copy of the malware to infect and control the targets of their choosing. This latter type of malware is known as a RAT or Remote Administration Trojan. In that case, hundreds or perhaps thousands of actors are each controlling their own small botnets using nearly identical code. In these types of Trojans, criminals use a "Builder Kit" to compile and customize their own malware to connect to infrastructure under their control. Examples of these kits include early versions of Zeus, SpyEye, Poison Ivy, DarkComet, Cutwail, BlackEnergy and others (Bodmer, 2011). Because FakeAV

operates in a “Pay Per Install” (PPI) affiliate model, many criminals are encouraged to compete against one another to try to get more victims to install their software rather than a peer affiliate's software (Caballero, 2011 & Cova, 2010) Since all the PPI affiliates spreading the same FakeAV are using the same malware, another method is needed to determine which malware is being spread by which affiliate. The most successful affiliates are those that have the most flexible or most enduring network infrastructure. In the related pharmaceutical affiliate programs, consistent infrastructure allows the top affiliates to earn over \$1M per year while the median affiliate receives \$3,000 or less (McCoy, 2012).

Both Winwebsec and FakeRean regularly change the user interface to better match the legitimate software that they are trying to emulate. In order for law enforcement to know whether a previously unseen version of the malware is actually part of a group they are investigating, some additional means of determining whether it is related or not may prove useful. Investigators often desire to identify affiliates as named individuals during the course of the investigation. “Turning” an affiliate can provide great insight into the management and practice of the overall criminal organization (Goodin, 2012). Through the techniques illustrated in this paper, law enforcement can sort out one FakeAV network from another and also identify the high value affiliate members that can further their investigations.

4. ENVIRONMENT

University of Alabama at Birmingham’s (UAB) malware analysis environment includes a Postgres database where malware statistics and metadata are stored, along with the actual binaries for more than 7 million malware samples dating back to May of 2008. As new malware is ingested, metadata is extracted and stored about the sample. The VirusTotal website is checked for the current detection of the sample by more than forty anti-virus providers and these results are stored (Canto, 2013). Most samples analyzed at UAB have been received in coordination with external parties, thus most of the samples have already been submitted to VirusTotal. Rather than submitting, the researchers at UAB search for the malware in the VirusTotal database by their MD5 hash and store the resulting vendor definitions for the malware in the “malware_detects” table.

As of January 10, 2014, the UAB Malware Repository contains 40,486 distinct malware samples that Microsoft identifies as Rogue:Win32/Winwebsec and an additional 13,231 distinct malware samples that Microsoft identifies as Rogue:Win32/FakeRean. While Microsoft has established that these are the two most dominant FakeAV families, most other vendors do not use this naming convention, and some vendors actually choose the opposite name for certain samples. For example, VIPRE antivirus labels 599 of the FakeRean samples as Winwebsec and 57 of the Winwebsec samples as FakeRean.

5. VENDOR AGREEMENT EXPERIMENT

The technique developed for building test sets of related malware uses a process called “*n*-way vendor agreement.”¹ In this technique, a script is passed parameters including a string that must be present in the malware name, a minimum number of vendors that must have used that string in their naming of a malware sample, and a date or date range in which the malware should have been initially reported. Although many vendors differ in their naming conventions, there are often certain “roots” found in many vendors naming choices. For example, many vendors use the phrase “zbot” to refer to all Zeus-related malware families. In this case, the term “fake” was used as the search term, and selected the fifty samples for each of seven days that had the most vendors who used the word “fake” to describe that malware. Although Microsoft prefers the prefix “Rogue”, many AV vendors use “FakeAV” or “FakeAlert” as their label for all families of FakeAV.

¹ The original “N-way Vendor Agreement” technique was developed with support from Sentar, Inc. in support of the DARPA Cyber Genome initiative.

One of the first tests was to determine whether the naming convention used by Microsoft was consistent with other vendors in the way it divided the world of FakeAV, and whether this naming convention would be useful while observing the network behavior of malware that had been identified as FakeAV. According to Microsoft’s SIR report, the two most common Fake AV families in the past year were Rogue:Win32/FakeRean and Rogue:Win32/Winwebsec. The researchers selected thirty samples of each from the database by asking for samples where vendor = ‘Microsoft’ and Malware_name = either ‘Rogue:W32/FakeRean’ or ‘Rogue:W32/Winwebsec’. A query was then performed to find which of the other anti-virus vendors used the same name as Microsoft for these 60 malware samples.

Because the researchers wanted to do a comparison across many anti-virus vendors, any of the 48 anti-virus products on the VirusTotal web site which did not detect at least 75% of these 60 samples were eliminated. The researchers also eliminated all but one from groups of vendors who labeled every sample the same as another vendor, indicating a shared detection engine. That left 18 vendors to consider. For each of these 18 vendors the following information was needed:

- a. How many different malware family names were assigned to these two Microsoft-labeled families?
- b. How many times was a single family name assigned to samples from BOTH Microsoft-labeled families?
- c. How many samples of FakeRean were not detected by this vendor?
- d. How many samples of Winwebsec were not detected by this vendor?
- e. How many samples were assigned a “generic” name?
- f. How many times did this vendor use the name FakeRean [and how many sub-variants did they assign]?
- g. How many times did this vendor use the name Winwebsec [and how many sub-variants did they assign]?

The eighteen vendors have been anonymized, A-R as shown in Table 1.

Table 1 Anonymized Anti-Virus Vendors

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R
a. # names	5	11	8	8	4	7	8	10	8	8	4	9	4	7	7	12	8	9
b. dupe name	1	1	2	1	1	0	1		1	0	0	1	1	0	1	1	1	0
c. missed FR	9	10	10	6	10	11	6	13	11	2	0	10	10	11	13	13	8	0
d. missed WWS	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
e. generic	0	22	25	16	37	24	16	2	27	0	2	16	0	17	9	0	18	24
f. FakeRean	0	2[2]	0	0	0	1	0	0	0	0	25	0	2	0	0	7	0	3 [2]
g. WWS	0	8[6]	0	0	0	0	0	1	0	0	0	0	0	0	0	5	0	10[2]

While Microsoft had only two names for these sixty malware samples, the average for these top performing vendors was 7.6 distinct family names used to name the samples.

Seventeen of the eighteen vendors detected 100% of the Winwebsec samples as malicious, but on the average, they failed to detect the FakeRean samples as being malicious 28% of the time (an average of 8.5 missed detections out of 30).

Fifteen of the eighteen vendors agreed that 2 of the samples were NOT FakeAV, but were rather Cosmu/KucIRC malware. All 18 labeled these two samples the same, regardless of the name they assigned.

Four vendors joined 2 Winwebsec samples with either 10, 11, or 13 FakeRean and labeled the group as “Kazy” malware. The same 2 Winwebsec samples were labeled the same way each time this occurred.

Given this sample, it seems that Winwebsec is highly detected and has high agreement on relationship; however, only four vendors used the name at all, and even the two most prominent only labeled 8 of 30 and 10 of 30 samples as Winwebsec. FakeRean has much less consistency and is more likely to be named outside the FakeAV family names. Only one vendor frequently used the FakeRean name (25 of 30 times) and only six vendors used it at all.

6. NETWORK EXPERIMENT AND ANALYSIS

UAB’s malware database was searched for widely detected malware samples that were believed to be FakeAV related to determine whether other characteristics, primarily demonstrated network traffic, could be used to cluster the malware into groups that would be useful to investigators. The supposition is that malware which connects to the common network infrastructure is certainly “related” to other malware samples that connect to the same infrastructure. In the paper *Driving in the Cloud*, researchers from Instituto Madrileño de Estudios Avanzados (IMDEA) perform a “milking” experiment where they repeatedly visit known drive-by infection sites to map over time what infections they are distributing, resulting in some very interesting insights into the common infrastructure relationships among malware families (Nappa, Rafique, & Caballero, 2013).

After an initial experiment performed by researchers at UAB of thirty-three FakeAV samples from the month of October 2013, it seemed that an experiment with a larger dataset was warranted.

Much work has already been done on categorizing malware families. Some of the common recent techniques have shown great ability to show the relationship between malware samples of the same family, including N-gram sequential pattern analysis (Liangboonprakong & Sornil, 2013), filtered block N-grams (Upchurch & Zhou, 2013), and highly parallel N-gram analysis (Jang, 2010). Others have used instruction frequency analysis (Han, Kang, & Im, 2011) and control flow graphs to detect malware families (Kang, 2011). Malware writers will change the source code of the malware in order to prevent detection. All of these techniques are extremely valuable, but they do not address the investigator’s question of how to tell if a single actor is using multiple (unrelated) malware families, or if multiple actors are using the same malware family in independent ways.

To answer these questions, the network connections made by executing each malware program were examined. The experiment was designed with a test set of 383 samples of FakeAV from the UAB Malware Repository. The query to perform this search, repeated for each day from November 4, 2013 to November 11, 2013 was:

```
./NwayVendorAgreement.sh -m fake -c 50 -d 2013-11-01
```

Combined with the initial thirty-three samples, this provided a dataset of 383 FakeAV samples to choose from. These were originally intended to be run through the Cuckoo Automated Sandbox environment (Guarnieri, Tanasi, Bremer, & Schloesser, 2014). However, a combination of problems arose. First, some of the AV has a very long “sleep” time before triggering which was problematic for the current automation environment. Secondly, some of the malware samples used anti-analysis techniques to determine they were in a virtual environment and failed to run at all (Chen, 2008). While the researchers look forward to building a more evasive Cuckoo environment (Ortega, 2012), for now it was decided to perform the work in a raw iron environment.

“Raw iron” in the malware analysis world refers to executing malware on native hardware and operating system, often with no analysis tools present on the execution environment (Kreibich, 2011). The test environment consisted of a raw iron machine on which the malware was executed and a machine that captured the network packets. The machines were connected to a hub, rather than a switch, for ease of packet eavesdropping from the Packet Capture machine as illustrated below.




Figure 1 Raw Iron Environment Network Diagram

Acronis True Image backup software was installed on the Malware Machine, and a backup to a separate drive was performed. A recovery boot disk was also created. The Try & Decide option of Acronis was used in order to refresh the Malware Machine after each run (Bayon, 2011). Originally intended to allow a consumer to install a piece of software and then “fully revert” to the pre-installed state or “commit to disk” if the change was acceptable, the application is perfect for malware analysis, allowing a raw iron environment to refresh the image in less than 90 seconds! Should the malware survive this recovery, the machine would be recovered using the recovery boot disk and the backup. Only one attempted sample failed to execute in this environment. Although still much faster, the lack of full automation caused the researchers to down-select the number of malware samples that were actually tested.

The first run, sub-selected from the October dataset, consisted of 18 FakeAV malware samples selected from the malware database using *n*-way vendor agreement. Acronis’ Try & Decide was started on the Malware Machine. The packet capture was started and the malware was launched. The malware was allowed to run while the network traffic was monitored with the packet capture software.




Figure 2 Sample FakeAV Scan

Once the FakeAV reached the point where its artificial scan was complete and the user was invited to “Remove all threats now” (Figure 2), the researchers connected to the purchase site, then stopped the packet capture, rebooted the Malware Machine, and used Acronis to discard all changes. This process was repeated for each of the 18 malware samples. Eleven distinct families were found based on the first end point contacted by the malware program, four of which had more than one member. When clustering on all network points, five distinct hosting clusters emerged.

Based on these positive results, the second batch of malware containing 350 samples was selected. Due to time considerations five groups of malware were chosen to be analyzed. The members of each of the five groups were selected based on proximity in size to the other group members. There were eleven 393k files, thirteen 396k files, twenty-eight 400k files, twenty-seven 401k – 404k files, and twenty-one 831k – 836k files.

For analysis, members of the two groups were combined together.

Data was first clustered based only on the first IP address contacted by the malware sample. That clustering produced the following result:

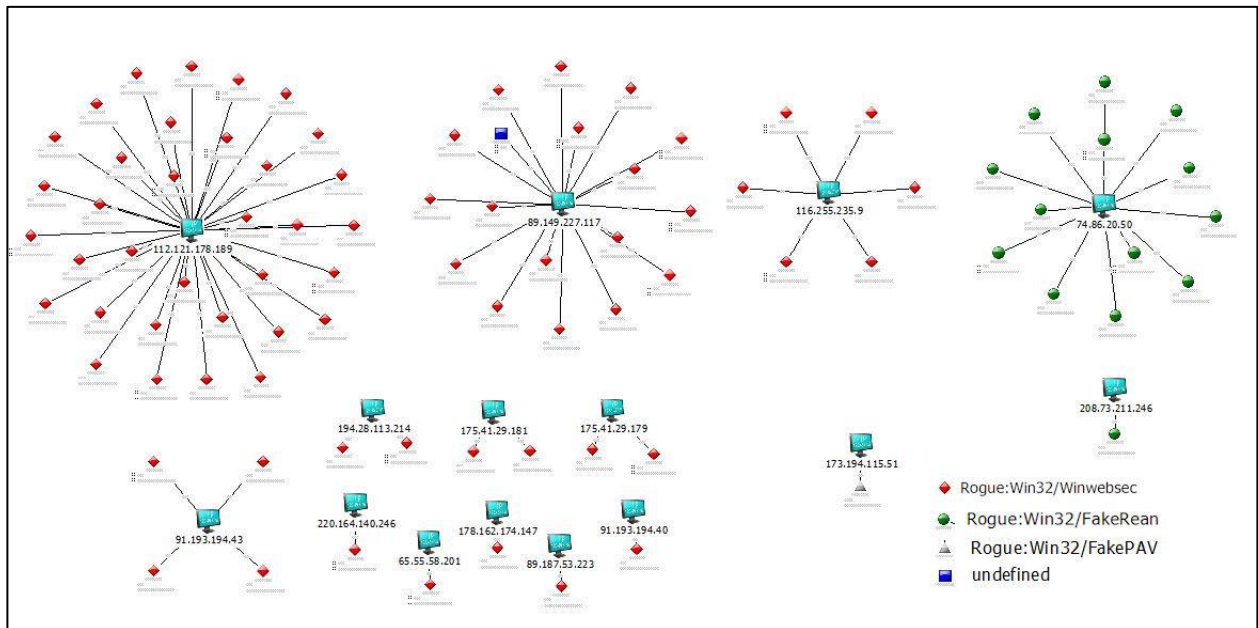


Figure 3 Clusters on First IP Address Contacted

In Figure 3, the malware has been color coded based on the definition assigned to it by Microsoft antivirus. If the top row clusters are labeled from left to right:

Cluster A consists of 30 Winwebsec samples that first contacted 112.121.178.189.

Cluster B consists of 17 Winwebsec and 1 unknown malware sample that first contacted 89.149.227.117.

Cluster C consists of 6 Winwebsec samples that first contacted 116.255.235.9.

Cluster D consists of 13 FakeRear samples that first contacted 74.86.20.50.

The remainder of the samples included a foursome, three pairs, and five singles that were Winwebsec samples, and two singleton clusters, one FakeRear, and the other “Rogue:Win32/FakePAV”.

Next, the malware was clustered with ALL network IP addresses contacted, further assisting in forming clusters. Rather than having the cluster nodes represent the malware family name, it was chosen to have the cluster nodes identify themselves by the TEMPLATE they portrayed:

Cluster A remained unchanged by this behavior, retaining the same number of samples and adding two IP addresses, with all samples identified as “Live Security Platinum”.

Cluster B expanded slightly to include 19 malware samples and 9 IP addresses, with all samples labeled as “Smart Fortress 2012”.

Cluster C retained the same number of samples (6) but now has 22 IP addresses, with five IP addresses being prominently linked between samples. Those were: 116.255.235.9, 141.8.224.79, 208.91.196.4, 184.51.150.146, and 69.43.161.163. Cluster C is also entirely composed of “Live Security Platinum” malware samples.




Figure 4 Two Live Security Platinum Clusters Show No Shared Infrastructure

Because “Live Security Platinum” is a Pay Per Install affiliate program, the diagram in Figure 4 demonstrates exactly the behavior one would expect to see if two (three if the singleton on the bottom right is counted) affiliates of the program are running competing infrastructure to drive their infections. Cluster A, on the top, has chosen to host on a long-term criminal IP address safely housed on Bullet-Proof Servers in China (Villeneuve, 2011). Cluster C, on the bottom, is using a wide variety of IP addresses scattered across many geographies and keeps in contact through the use of redundant IP addresses, so that if part of the affiliate’s diversified infrastructure is disrupted, they can still connect via backup IP addresses. Both are pushing the same malware family, but the infrastructure makes it clear there are two (or more) separate affiliates involved.

Cluster D remained unchanged, with all nodes identified as “Internet Security”.

More interesting was that the additional IP addresses being added to the mix created two new types of clusters that are named Cluster E (Figure 5) and Cluster F (Figure 6).

Cluster E daisy chains together from top to bottom:

- *XP Anti Spyware 2011* to *XP Security 2011* by the common IP 208.73.211.246
- 216.166.16.134 connects a cluster of five *MS Removal Tool* samples which share the common IPs 69.50.195.76, 69.50.209.220, and 91.193.194.43
- 69.50.195.76 joins a cluster of three *System Tool* samples to the common IPs 194.28.113.214 and 212.71.10.110.

Cluster F joins together 5 *System Care Anti Virus* samples and 15 IP addresses.




Figure 5 Cluster E - Mixed Care Anti-Virus




Figure 6 Cluster F – System

7. CONCLUSION

Law enforcement and other malware investigators need to be able to consistently determine whether a given complaint or malware sample is related to their investigation or not. The researchers have demonstrated that reliance on malware names assigned by a single Anti-virus vendor is inadequate to this task. By selecting malware to analyze using the *n*-way vendor agreement technique, patterns of malware family names that are likely to be consistent across vendors can be learned. Through observation of groups of potentially related malware selected through this technique, the relationships of malware samples can be diagramed by their shared infrastructure. The network relationships documented in this way can be used to cluster not only malware related by common code, but to assign attribution of those who control and distribute malware based on their hosting decisions.

While it is always true that some samples are missed by some AV vendors, and there are still many uses of “generic” naming terms, a combination of cross-vendor name agreements, compared with analysis of the network infrastructure addressed by the malware, can provide deep insights into the relationships between malware samples and the way that the same malware differs based on hosting decisions made by the personnel behind the malware distribution.

REFERENCES

- Antonio Nappa, M. Z. (2013). Driving in the Cloud: An analysis of drive-by download operations and abuse reporting. In P. S.-P. Konrad Rieck, *Detection of Intrusions and Malware, and Vulnerability Assessment*, 1-20. SpringerLink.
- APWG. (2006). The crimeware landscape: Malware, phishing, identity theft and beyond. Retrieved on January 9, 2014, from Anti-Phishing Working Group http://docs.apwg.org/reports/APWG_CrimewareReport.pdf
- Bayon, D. (2011). Acronis true image home 2012 review. Retrieved on January 12, 2014 from PC Pro <http://www.pcpro.co.uk/reviews/software/370153/acronis-true-image-home-2012>
- Bodmer, S. (2011). It's raining source. Retrieved on January 9, 2014 from Damballa Blog: The Day Before Zero <https://blog.damballa.com/archives/1313>
- Brett Stone-Gross, M. C. (2009). *Your Botnet is My Botnet*. CCS '09, 635-647. New York, NY: ACM.
- Caballero, J. G. (2011). Measuring pay-per-install: The commodotization of malware distribution. Usenix security symposium.
- Canto, J. (2013). About VirusTotal. Retrieved on January 12, 2014 from VirusTotal.com <https://www.virustotal.com/en/about/>
- CERT Polska. (2013). Technical report: Zeus-P2P monitoring and analysis. Retrieved on January 10, 2014 from CERT Polska http://www.cert.pl/PDF/2013-06-p2p-rap_en.pdf
- Chen, X. A. (2008). Towards an understanding of anti-virtualization and anti-debugging behavior in modern malware. *IEEE International Conference on Dependable Systems and Networks*, 177-186.
- Claudio Guarnieri, A. T. (2014). Automated malware analysis. Retrieved on January 12, 2014 from Cuckoo Sandbox <http://www.cuckoosandbox.org/about.html>
- Cova, M. L. (2010). An analysis of rogue AV campaigns. *Recent Advances in Intrusion Detection (RAID '10)*, 442-463. Springer Berlin Heidelberg.
- FBI Press. (2011). Department of Justice disrupts international cyber crime rings distributing scareware. Retrieved on December 29, 2013 from FBI National Press Releases <http://www.fbi.gov/news/pressrel/press-releases/department-of-justice-disrupts-international-cybercrime-rings-distributing-scareware>
- Federal Trade Commission v. Innovative Marketing, Inc., 08-CV-3233-RDB Federal Court District of Maryland, December 10, 2008.
- FTC. (2013). Innovative Marketing, Inc., et al. Retrieved on January 9, 2014 from FTC Cases and Proceedings <http://www.ftc.gov/news-events/press-releases/2012/10/ftc-case-results-163-million-judgment-against-scareware-marketer>
- Goodin, D. (2012). Turncoat hackers: A brief history of snitching in high-tech dragnets. Retrieved on January 10, 2014 from Ars Technica <http://arstechnica.com/business/2012/03/turncoat-hackers-a-history-of-snitching-in-high-tech-dragnets/>
- Han, K. S., Kang, B., & Im, E. G. (2011). Malware classification using instruction frequencies. 2011 ACM Symposium on Research in Applied Computation, 298-300. New York, NY: 2011.
- Jang, J. D. (2010). Bitshred: Fast, scalable malware triage. Pittsburgh, PA: Cylab, Carnegie Mellon University.
- John, J. P., Yu, F., Xie, Y., & Abadi, M. (2011). deSEO: Combating search-result poisoning. USENIX Security Symposium.

- K, S. (2011, June 19). Gagarincash AV Affiliate. Retrieved on January 20, 2014 from XyliBox: Tracking Cyber Crime <http://www.xylibox.com/2011/06/tracking-cyber-crime-gagarincash-av.html>
- Kang, B. K. (2011). Fast malware family detection method using control flow graphs. RACS '11 Proceedings to the 2011 ACM Symposium on Research in Applied Computation, 287-292. ACM.
- Krebs, B. (2011). Fake Antivirus Down, But Not Out. Retrieved on January 10, 2014 from Krebs On Security <http://krebsonsecurity.com/2011/08/fake-antivirus-industry-down-but-not-out/>
- Kreibich, C. W. (2011). GQ: Practical containment for measuring modern malware systems. Proceedings of the 2011 ACM SIGCOMM conference on Internet Measurement, 397-412. ACM.
- Liangboonprakong, C., & Sornil, O. (2013). Classification of malware families based on n-grams sequential pattern features. 8th IEEE Conference on Industrial Electronics and Applications (ICIEA), IEEE, 777-782. Melbourne.
- McCoy, D. P. (2012). PharmaLeaks: Understanding the business of online pharmaceutical affiliate programs. USENIX Security Symposium.
- Michael Bailey, J. O. (2007). Automated classification and analysis of internet malware. RAID '07 Proceedings of the 10th international conference on Recent Advances in Intrusion Detection, 178-197. Berlin: Springer-Verlag.
- Microsoft. (2013). Microsoft Security Intelligence Report Volume 15. Redmond, OR: Microsoft.
- Mimoso, M. (2014). Malicious ads on DailyMotion redirect to fake AV attack. Retrieved on January 9, 2014 from ThreatPost <http://threatpost.com/malicious-ads-on-dailymotion-redirect-to-fake-av-attack/103494>
- Ortega, A. (2012). Hardening cuckoo sandbox against VM aware malware. Retrieved on January 10, 2014 from AlienVault <http://www.alienvault.com/open-threat-exchange/blog/hardening-cuckoo-sandbox-against-vm-aware-malware>
- Provos, N., Mavrommatis, P., Rajab, M. A., & Morose, F. (2008). All your iFRAMEs point to us. USENIX Security Symposium.
- Rascagnères, P. (2013). APT1: Technical Backstage. Retrieved on January 11, 2014 from itrust consulting http://www.malware.lu/Pro/RAP002_APT1_Technical_backstage.1.0.pdf
- Samosseiko, D. (2009). The Partnerka-what is it, and why should you care? Virus Bulletin Conference, 115-120.
- Stone-Gross, B., Abman, R., Kemmerer, R., Kruegel, C., Steigerwald, D., & Vigna, G. (2013). The Underground Economy of Fake Antivirus Software. Economics of Information Security and Privacy III, 55-78.
- Upchurch, J., & Zhou, X. (2013). First byte: Force-based clustering of filtered block N-grams to detect code reuse in malicious software. 8th International Conference on Malicious and Unwanted Software, IEEE, 68-76. Fajardo, PR, USA.
- Villeneuve, N. (2011). Targeting the source: FakeAV affiliate networks. Retrieved on January 2014 from Trend Micro <http://www.trendmicro.com/media/wp/fakeav-affiliate-networks-whitepaper-en.pdf>
- Warner, G. (2008). FTC moves against fake AntiVirus "ScareWare" companies. Retrieved on January 7, 2014 from CyberCrime & Doing Time <http://garwarner.blogspot.com/2008/12/ftc-moves-against-fake-av-scareware.html>

