Abstract
Advanced Persistent Threat (APT) attacks are highly organized, continue for prolonged periods, and exhibit discernible attributes or patterns. To maintain command and control network redundancy, APT attacks are generally embedded with multiple domain name system (DNS) names. Intuitively, APT attackers maintain and control a large number of DNS–internet protocol (IP) address pairs. The majority of existing studies on malware attribution greatly emphasize grouping the technological or behavioral contexts from malware binaries. In this study, we examined a small sample of malware from a specific victim group that had been subjected to APT attacks. Results show that the attackers followed specific behavioral patterns involving registering DNS domains and frequent use of stable DNS–IP pairs. Although gathering such evidence from malware binaries is not complicated, tedious online queries regarding open source information are required. We developed an automated solution to simplify the collection and storage of information as a knowledge base for future analysis. Once the initial set of malicious DNS–IP address pairs, "parked domains," and "whois information" are identified, the database can be used to perform updates manually. This database can be used for further analysis using a visualization tool as well as to identify the possible identities or personas of the attackers. We used Maltego to perform the analysis in this study.

1. Introduction

The attribution of malware attackers pose a challenge to malware analysts. Most researchers extract technological artifacts from malware binaries and perform data mining analysis to determine the identity of attackers or at least fingerprint partial information of malware authors. When studying Windows malware, the portable executable (PE) headers are first “deconstructed.” The extracted PE metadata are then categorized in accordance with a set of defined rules, stored in an SQL database, and further analyzed (Yonts, 2012). A number of malware analysts extend their work to contextual analysis by obtaining attributes or “genes” from different “layers,” such as the exploits or shell code used, the metadata of the PE information, the connected Transmission Control Protocol port number, and the command and control (C2) network infrastructure (Xecure-Lab, 2012). Meanwhile, other researchers take a step backward to extract the metadata information contained in email headers if the malware was distributed through spear-phishing emails (Lee, M.& Lewis, D., 2011). Some malware analysts group attackers by referring to proprietary reverse engineering and behavioral analysis technology (Digital DNA, 2014). All of these researchers claim to have successfully identified various attacker groups, which are sometimes given a code name, such as APT1 (Mandiant), Comment Crews (Hoglund, G.), Soysauce (HB Gary), or DeepPanda (CrowdStrike). The researchers have given hints on how they categorized the groups; however, none of them, except for Mandiant’s APT1 Report, have provided the complete details of their work.

Since Mandiant published the APT1 Report, a considerable number of researchers have studied the C2 infrastructure of attackers in a similar manner. They use virtualization tools such as Maltego to illustrate their findings on how

to associate these attacker groups (HBGary, 2012). Maltego graphs are useful to analysts for illustration purposes; however, the information displayed in these graphs is only a representation of static analytical results at the time of the queries in question. Our study begins in a similar manner, but we place more emphasis on monitoring the changes in network infrastructure or domain name usage over a period of time, say one year. We have developed an automated solution to simplify the tasks of gathering and storing relevant information, and which can be used as a knowledge base for future analysis. Once the initial set of malicious domain name system–internet protocol (DNS–IP) address pairs, "parked domains," and "whois information" are identified, the database can be used to perform updates manually. This database can be used for further analysis using a visualization tool as well as to identify the possible identities or personas of attackers. In our study, we used Maltego to perform the analysis.

To maintain C2 network redundancy, Advanced Persistent Threat (APT) attacks are generally embedded with multiple DNS names (two-level domains) in the attacking weapons. An intuitive view is that APT attackers maintain and control large number of DNS–IP address pairs. We believe that the attackers might assign a team to register large numbers of DNS domain names for such purposes. For each fresh registration, they leave unique and identifiable personal information on the “whois” servers. The “whois” servers may provide reliable links, particularly at the stage between registration and first launch of the malware.

In this study, we examine two aspects of the network infrastructure: (1) the DNS–IP address pairs and (2) the “whois information.” We believe that the information contains the highest possible intelligence that links the malware to the human actors.

2. Related research

During the 2010 Blackhat event, Hoglund presented a paper titled Malware Attribution: Tracking Cyber Spies & Digital Criminals (Hoglund, 2010). He argued that Social Cyberspace (i.e., DIGINT) and Physical Surveillance (i.e., HUMINT) were the keys with which to identify the actors behind malicious software attacks. However, Hoglund also pointed out the near impossibility of finding the human actors with definitive intelligence, such as the social security numbers or physical locations of the attackers. However, he described forensics marks that could be extracted from raw data in three intelligence layers: (a) Net Recon (C2), (b) Developer Fingerprints, and (c) Tactics, Techniques, and Procedures (TTP). Among these three layers, TTP should carry the highest intelligence value for identifying human attackers.

Basing on Hoglund’s theory, Boman developed a tool called VXCage, which extracts technical metadata from binaries and store the artifacts in a relational database for further analysis (Boman, 2013).

Pfeffer et al. (2012) investigated how to extract “genetic information” from malware by reverse engineering critical PE header information and following

evolutionary traces and functional linguistics from binaries. However, their research does not provide direct links to DIGINIT or HUMINT at the higher end of the Intelligence Spectrum as described by Hoglund in his Blackhat presentation.

By contrast, our study tracks changes in DNS domains and "whois information" not only during the first instance of being found but for a long period after the changes have been identified. Our analysis confirms that the "whois information" retained in the DNS domain network infrastructure remains more stable than the fast flux infrastructure used by publicly distributed malware. By monitoring the network infrastructure and "whois information" of "parked domains" from the malicious IP addresses of APT attacks, the resulting knowledge base may provide more relevant direct links to the DIGINIT or HUMINT of the attackers.

3. APT attack attributes and patterns

3.1 APT life cycle

The majority of the existing literature describe the APT cycle by means of a series of processes: (a) defining the target; (b) sending spear-phishing; (c) dropped backdoor; (d) initial outbound connection; (e) data gathering; and (f) exfiltration (Fig. 1).

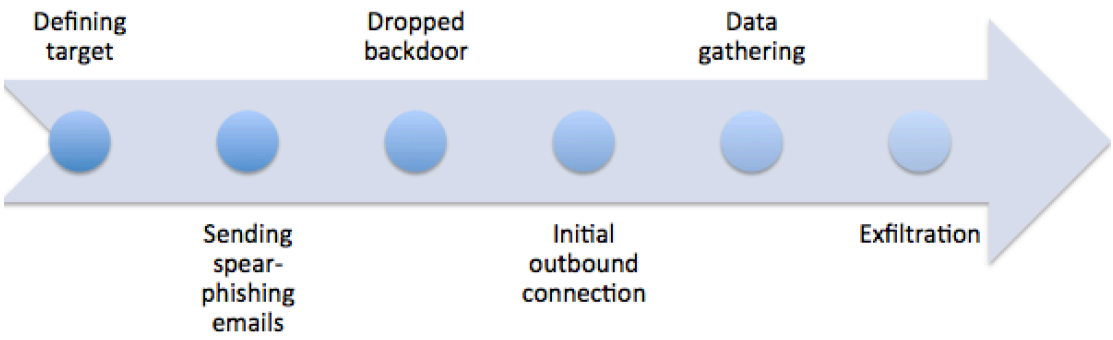

APT life cycle (Fig. 1)

We suggest that the life cycle should be extended further to become initiated by means of social engineering of the victims (described as "human reconnaissance" for the purpose of preparing spear-phishing emails) and expanded to lateral movements (pushing more tools or collecting selective intelligence from victims' machines) after the initial exfiltration (Fig. 2). These processes require increased human interaction, which may provide direct links to the DIGINIT or HUMINT.

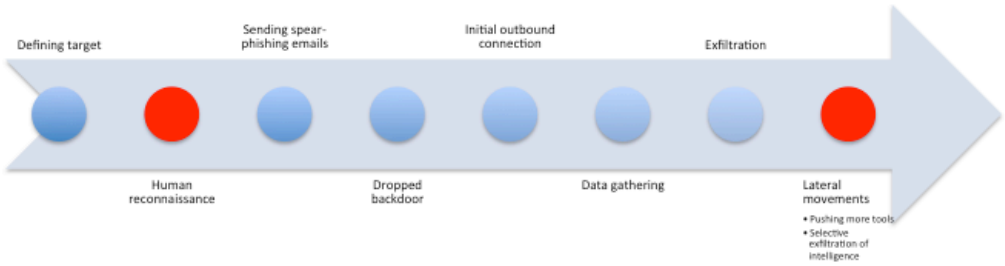

Extended APT life cycle (Fig. 2)

3.2 APT infrastructure tactics

APT attacks are highly organized and are launched by a group of attackers over a prolonged period. APT attacks exhibit discernible attributes or patterns. APT attackers tend to register and control their DNS domain network infrastructures in the following workflow:

- Domain registration
- Naming of domains: domain naming convention is not typo squatting, but follows a pattern of meaningful Chinese PingYing （拼音）
- Creation of second-level domain name and IP address pairs
- Engaging a “friendly ISP” to use a portion of their C-class subnet of IP addresses situated at the domicile of the targeted victims (for the purpose of evasion of blacklisting or to make the communication geologically viable)
- Reusing DNS names and IP addresses: DNS names and IP addresses may be cycled for reuse (a.k.a. campaigns), which may provide indications or links to the attacker groups
- Embedding multiple DNS A-records in exploits
- Preparing spear-phishing email content after reconnaissance of the targeted victims
- Launching malicious attachments through spear-phishing emails
- Collection of intelligence: the exploits drop binaries that extract the DNS records and begin communicating with the C2 by resolving the IP

addresses from DNS servers. The C2 servers or C2 proxies register the infectioned machine on the C2 database and the intelligence analysts of the attacker groups review the preliminary collected information of the targeted victims through C2 portals. The infected machines are further instructed to perform exfiltration of collect further intelligence from the infected machines.

- Manipulation of domains: the infrastructure technical persons of the attacker group apply changes (domain manipulation) to the DNS-IP address pair, domain name registration information (Whois information), and the "parked domains" from time to time or when a specific incident occurs (such as the takedown of the C2 proxies by law enforcement or being identified as bogus by local CERT)
- Frequency of change of information: In contrast with the Fast-Flux Services Networks mentioned by the HoneyNet Project, the information does not change with high frequency
- Monitoring of DNS–IP Address pairs and Whois information: as a result of the domain manipulation activities, the content of the DNS-IP Address pair and Whois information should be monitored immediately after being identified

3.3 Role of DNS in APT Attacks

When a registrant wants to set up a DNS domain, he or she is required to provide the domain name registrars with basic information, such as valid email address, name, street address, and NS records pointing to authoritative name servers. Although this information can be changed after registration, we assume that the APT attackers have registered a large number of the DNS domains and are too lazy to keep changing this information immediately.

Once the DNS domain is registered, in most cases, the registrant may use the domain name to create an A record and "park" the domain name to a controlling IP address. When an attack campaign is launched, the registrant creates more DNS–IP address pairs for their weapons by adding "A records" to their authoritative name servers. These newly created DNS–IP address pairs are then embedded into their weapons and sent out with spear-phishing emails as attachments.

An intuitive view is that APT attackers maintain and control a large number of DNS–IP address pairs and keep the "whois information" unchanged during the early stage after its first use. To characterize the DNS domain registration pattern, we collect the intelligence from an open source. Owing to the imperfection of the IP, the format of the information provided by the "whois" servers does not follow a definitive format. At present, no perfect automated solution exist for collecting the "whois information" of these "parked" domains. In addition, no database keeps all zone record changes for every DNS domain and "whois information." Although the Passive DNS Project offers a partial solution for querying historical records of DNS and IP address pairs, a complete knowledge base can still be provided.

3.4 VirusTotal and Passive DNS Project

According to Virus Total, Passive DNS is a technology that constructs zone replicas without cooperation from zone administrators, and is based on captured name server responses. Passive DNS is a highly scalable network design that stores and indexes both historical DNS data that can help answer questions such as: (a) where did this domain name point to in the past and (b) which domain name points to a given IP network. Virus Total scans thousands of samples per day; these samples communicate with certain domains or IP addresses and may therefore have a privileged position to gather highly malicious chance passive DNS information.

4. Our case findings

Our original research objective was to examine DNS–IP address pairs and "whois information" changes for malicious domains to identify attacker groups. However, malware attribution is not a straightforward task. Our approach required a high volume of technical information and contextual analysis from samples collected in their early stage of distribution. Building up a relevant knowledge base requires more live samples and tedious online queries for open source information, as described above. During our research period, we studied a small sample set: 28 samples linked to 61 IP addresses, 370 "parked" DNS domains, and 922 "whois" records. Our study shows that attackers follow specific behavioral patterns when registering DNS domains and frequently use stable DNS–IP pairs. Unfortunately, our hypothesis requires a significantly larger amount of DNS–IP address pairs from live samples to generate observable statistics to concretely identify who the attackers are.

The initial empirical findings of most of our samples are discussed in Ran2's blog (Ran2, 2013). Although gathering such evidence from malware binaries is not complicated, tedious online queries for open source information are required. We have developed an automated solution to simplify the tasks of collecting and storing the information as a knowledge base for further analysis.

Basing on this sample set and using our automated solution, we captured more DNS-IP address pairs and "whois information." Even though most of the DNS-IP addresses and "whois information" were modified, we performed attribution analysis based on the collected data and used Maltego to depict the results (Fig. 3). An interactive analysis work was uploaded to Youtube.

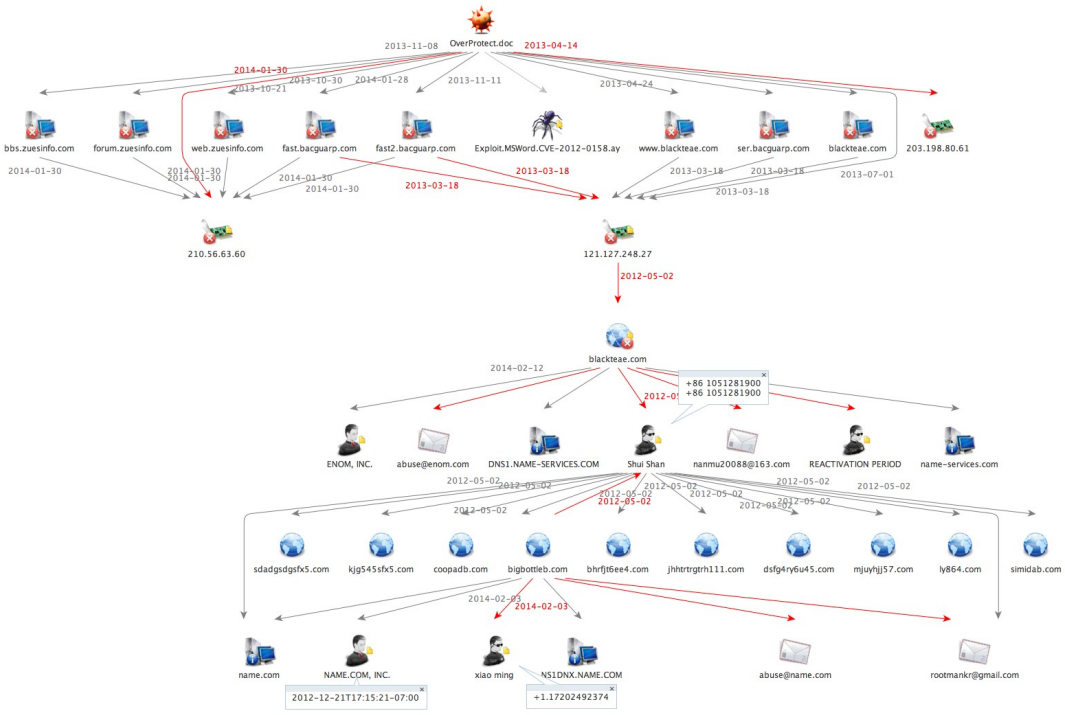

Clustering of network infrastructure for the OverProtect sample (Fig. 3)

5. Further research

We are going to publish our automated solution (Ran2 et al., 2014) after we have polished the source code. We believe that by using the proposed solution, any independent researcher can perform attribution analysis and further update the solution through plugins by extracting more intelligence artifacts to update their knowledge base.

We understand that identifying the human actors cannot completely prevent APT attacks. However, we believe that by monitoring the network infrastructure for definitive malicious DNS–IP address pairs, we may be able to generate a blacklist that can identify potential attacks at an early stage.

6. Conclusion

In this paper, we have discussed our intuitive views on the TTP of APT attackers. By continuously monitoring “whois servers” and DNS–IP address pairs, we may be able to collect more traces with which to identify the attackers. Gathering such intelligence is not complicated, but tedious online queries for open source information are necessary. We have developed an automated solution to simplify the tasks of collecting and storing the information as a knowledge base for future analysis. Our case findings demonstrate that APT attackers follow certain behavioral patterns when launching their past attacks. However, intelligence may be lost if they change their TTP in the future, particularly after the publication of this paper. Lastly, we believe that TTP are determined by the

cultural background of the attacker groups; the intelligence collection process should thus be adjusted toward these changes and analysts should have the same cultural mindset to identify meaningful evidence that would help identify the attackers.

7. Acknowledgement

This work was supported by the VX Research Group with thanks to Kenneth Tse, Frank Ng, Roland Cheung and Anthony Lai.

Reference

APT1 Report (2012), Retrieved on 9 April 2014, from Mandiant: http://intelreport.mandiant.com/Mandiant_APT1_Report.pdf

Boman, M. (2013). Malware Dataming and Attribution, Retrieved on 9 April 2014, from DeepSec 2013 Talk: http://blog.deepsec.net/?p=1697

Digital DNA, Retrieved on 9 April 2014, from HBGary: http://hbgary.com/products/digital_dna™

HBGary (2012). APT World At War, Retrieved on 9 April 2014, from HBGary: https://hbgary.com/sites/default/files/publications/APTWORLDATWARCHINA1600x1200.pdf

Hoglund, G. (2010). Malware Attribution: Tracking Cyber Spies & Digital Criminals. Retrieved on 9 April 2014, from YouTube: https://www.youtube.com/watch?v=Fch0bQ5UPmM

Lee, M., & Lewis D. (2011) Clustering Disparate Attacks: Mapping The Activities of the Advanced Persistent Threat. Retrieved on 9 April 2014, from VirusBulletin: https://www.virusbtn.com/pdf/conference_slides/2011/Lee-VB2011.pdf

Pfeffer, A. et al. (2012). Malware Analysis and attribution using Genetic Information, Proceedings of the 2012 7th International Conference on Malicious and Unwanted Software (MALWARE)

Ran2. (2013). Tracing APT163QQ. Retrieved on 9 April 2014, from espionageware's blog: http://espionageware.blogspot.hk/2013/05/tracing-apt163qq.html

Ran2 et al. (2014). Automated Solution. Retrieved on 1 July 2014, from Google Code: https://code.google.com/p/malicious-domain-profiling/

Xecure-Lab (2012). Prepare for the "Advanced Persistent Threat" Warefare, Retrieved on 9 April 2014, from Xecure-Lab: http://blog.xecure-lab.com/2012/07/prepare-for-advanced-persistent-threat.html

Yonts, J. (2012). Attributes of Malicious Files. Retrieved on 9 April 2014, from SANS Institute: http://www.sans.org/reading-room/whitepapers/malicious/attributes-malicious-files-33979

YouTube (2014). Retrieved on 9 April 2014, from Ran2: http://youtu.be/n3wnAcgjvog