Detecting and Tracking the Rise of DGA-Based Malware

MANOS ANTONAKAKIS, ROBERTO PERDISCI, NIKOLAOS VASILOGLOU, AND WENKE LEE

Manos Antonakakis received his PhD in computer science from Georgia Institute of Technology under Professor Wenke Lee's supervision. Currently, he works at Damballa as the Director of Academic Sciences where he is responsible for academic research, university collaborations, and technology transfer efforts. His main research interests are in network security and machine learning/data mining. manos@damballa.com

Roberto Perdisci is an Assistant Professor in the Computer Science Department at the University of Georgia, and an Adjunct Assistant Professor in the School of Computer Science at the Georgia Institute of Technology. He is the recipient of a 2012 NSF CAREER Award. His main research interests are in network security and machine learning/data mining. perdisci@cs.uga.edu

Nikolaos Vasiloglou received his PhD in electrical engineering from the Georgia Institute of Technology. He has extensive experience in developing machine-learning applications and algorithms. In the past he has developed machine-learning engines and models for several companies. nvasil@ieee.org

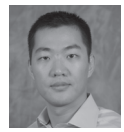

Wenke Lee is a Professor in the School of Computer Science at Georgia Institute of Technology and the Director of the Georgia Tech Information Security Center (GTISC). He earned his PhD in computer science from Columbia University in 1999. He has published more than 100 scholarly articles. His current research projects are in the areas of botnet detection, malware analysis, virtual machine monitoring, and Web 2.0 security and privacy, with funding from NSF, DHS, DoD, and industry. wenke@cc.gatech.edu

When bots go in search of their command and control (C&C) servers, they often use algorithmically generated domain names (DGAs). We have created a system (Pleiades) that watches unsuccessful DNS resolution requests (NXDomain) from recursive DNS servers in large networks. Pleiades can reliably identify new clusters of NXDomains generated by DGAs, the newly infected hosts, and often, the actual C&C servers the DGA malware employs. In this article, we explain how our system works, as well as the most interesting information about current bot infections and C&C structures.

Introduction

Botnets are groups of malware-compromised machines, or bots, that can be remotely controlled by an attacker (the botmaster) through a command and control (C&C) communication channel. Botnets have become the main platform for cybercriminals to send spam, steal private information, host phishing Web pages, etc. Over time, attackers have developed C&C channels with different network structures. Most botnets today rely on a centralized C&C server, whereby bots query a predefined C&C domain name that resolves to the IP address of the C&C server from which commands will be received. Such centralized C&C structures suffer from the "single point of failure" problem because if the C&C domain is identified and taken down, the botmaster loses control over the entire botnet.

To overcome this limitation, attackers have used P2P-based C&C structures in botnets such as Nugache, Storm, and more recently, Waledac, Zeus, and Alureon (aka TDL4). Whereas P2P botnets provide a more robust C&C structure that is difficult to detect and take down, they are typically harder to implement and maintain. In an effort to combine the simplicity of centralized C&Cs with the robustness of P2P-based structures, attackers have recently developed a number of botnets that locate their C&C server through automatically generated pseudo-random domains names. In order to contact the botmaster, each bot periodically executes a domain generation algorithm (DGA) that, given a random seed (e.g., the

current date), produces a list of candidate C&C domains. The bot then attempts to resolve these domain names by sending DNS queries until one of the domains resolves to the IP address of a C&C server. This strategy provides a remarkable level of agility because even if one or more C&C domain names or IP addresses are identified and taken down, the bots will eventually get the IP address of the relocated C&C server via DNS queries to the next set of automatically generated domains. Notable examples of DGA-based botnets (or DGA-bots, for short) are Bobax, Kraken, Sinowal (aka Torpig), Srizbi, Conficker-A/B/C, and Murofet.

A defender can attempt to reverse engineer the bot malware, particularly its DGA algorithm, to pre-compute current and future candidate C&C domains in order to detect, block, and even take down the botnet; however, reverse engineering is not always feasible because the bot malware can be updated very quickly (e.g., hourly) and obfuscated (e.g., encrypted, and only decrypted and executed by external triggers such as time).

We recently proposed a novel detection system, called Pleiades [1], capable of identifying DGA-bots within a monitored network without reverse engineering the bot malware. Pleiades is placed between the network machines and the local recursive DNS (RDNS) server (aka "below" the recursive DNS infrastructure of the network) or simply at the edge of a network to monitor DNS query/response messages from/to the machines within the network. Specifically, Pleiades analyzes DNS queries for domain names that result in Name Error responses, also called "NXDOMAIN" responses, i.e., domain names for which no IP addresses (or other resource records) exist.

The focus on NXDomains is motivated by the fact that modern DGA-bots tend to query large sets of domain names among which relatively few successfully resolve to the IP address of the C&C server. Therefore, to identify DGA domain names automatically, Pleiades searches for relatively large clusters of NXDomains that (1) have similar syntactic features and (2) are queried by multiple potentially compromised machines during a given epoch.

The intuition is that in a large network, such as the ISP network where we ran our experiments, multiple hosts may be compromised with the same DGA-bots. Therefore, each of these compromised assets will generate several DNS queries resulting in NXDomains, and a subset of these NXDomains will likely be queried by more than one compromised machine. Pleiades automatically is able to identify and filter out "accidental" user-generated NXDomains due to typos or misconfigurations. When Pleiades finds a cluster of NXDomains, it applies statistical learning techniques to build a model of the DGA. This is used later to detect future compromised machines running the same DGA and to detect active domain names that "look similar" to NXDomains resulting from the DGA and therefore probably point to the botnet C&C server's address.

Overview of Pleiades

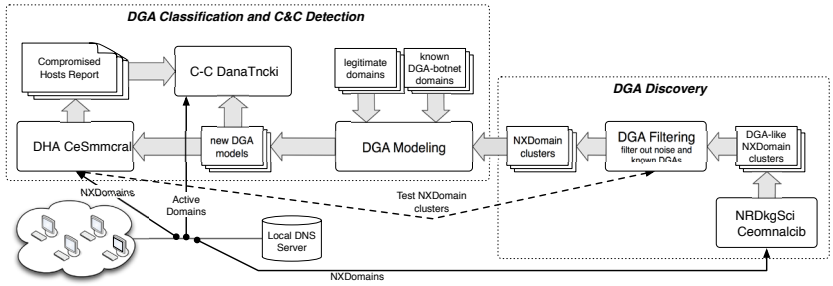

Figure 1: A high-level overview of Pleiades

Next, we provide a high-level overview of our DGA-bot detection system, Pleiades. As shown in Figure 1, Pleiades consists of two main modules: a DGA Discovery module and a DGA Classification and C&C Detection module. We discuss the roles of these two main modules and their components, and how they are used in coordination to learn actively and update DGA-bot detection models.

DGA Discovery

The DGA Discovery module analyzes streams of unsuccessful DNS resolutions, as seen from "below" a local DNS server (see Figure 1). All NXDomains generated by network users are collected during a given epoch (e.g., one day). Then, the collected NXDomains are clustered according to the following two similarity criteria: (1) the domain name strings have similar statistical characteristics (e.g., similar length, level of "randomness," character frequency distribution, etc.), and (2) the domains have been queried by overlapping sets of hosts. The main objective of this NXDomain clustering process is to group together domain names that likely are automatically generated by the same algorithm running on multiple machines within the monitored network.

Naturally, because this clustering step is unsupervised, some of the output NXDomain clusters may contain groups of domains that happen to be similar by chance (e.g., NXDomains due to common typos or to misconfigured applications). Therefore, we apply a subsequent filtering step. We use a supervised DGA Classifier to prune NXDomain clusters that appear to be generated by DGAs that we have previously discovered and modeled, or that contain domain names that are similar to popular legitimate domains. The final output of the DGA Discovery module is a set of NXDomain clusters, each of which likely represents the NXDomains generated by previously unknown or not yet modeled DGA-bots.

DGA Classification and C&C Detection

Every time a new DGA is discovered, we use a supervised learning approach to build models of what the domains generated by this new DGA "look like." In particular, we build two different statistical models: (1) a statistical multi-class classifier that focuses on assigning a specific DGA label (e.g., DGA-Conficker.C) to the set of NXDomains generated by a host h_i and (2) a hidden Markov model (HMM) that focuses on finding single active domain names queried by h_i that are likely

generated by a DGA (e.g., DGA-Conficker.C) running on the host, and are therefore good candidate C&C domains.

The DGA Modeling component receives different sets of domains labeled as Legitimate (i.e., "non-DGA"), DGA-Bobax, DGA-Torpig/Sinowal, DGA-Conficker.C, New-DGA-v1, New-DGA-v2, etc., and performs the training of the multi-class DGA Classifier and the HMM-based C&C Detection module.

The DGA Classification module works as follows. Similar to the DGA Discovery module, we monitor the stream of NXDomains generated by each client machine "below" the local recursive DNS server. Given a subset of NXDomains generated by a machine, we extract a number of statistical features related to the NXDomain strings. Then we ask the DGA Classifier to identify whether this subset of NXDomains resembles the NXDomains generated by previously discovered DGAs. That is, the classifier will either label the subset of NXDomains as generated by a known DGA, or tell us that it does not fit any model. If the subset of NXDomains is assigned a specific DGA label (e.g., DGA-Conficker.C), the host that generated the NXDomains is deemed to be compromised by the related DGA-bot.

Once we obtain the list of machines that appear to be compromised with DGA-based bots, we take the detection one step further. While all previous steps focused on NXDomains, we now turn our attention to domain names for which we observe valid resolutions. Our goal is to identify which domain names, among the ones generated by the discovered DGA-based bots, actually resolve into a valid IP address. In other words, we aim to identify the botnet's active C&C server.

To achieve this goal, we consider all domain names that are successfully resolved by hosts that have been classified as running a given DGA, say New-DGA-vX, by the DGA Classifier. Then we test these successfully resolved domains against an HMM specifically trained to recognize domains generated by New-DGA-vX. The HMM analyzes the sequence of characters that compose a domain name `d`, and computes the likelihood that `d` is generated by New-DGA-vX.

DGA Discoveries and Case Studies

In this section, we present the most important experimental results of our system. We will elaborate on the DGAs we discovered throughout the two years of NXDomain monitoring period at a large US ISP. Then we will summarize the most interesting findings from the 13 DGAs we detected. Seven of them use a DGA algorithm from a known malware family. The other six, at the time of discovery and to the best of our knowledge, have no known malware association. We will conclude with three cases studies of currently active threats that employ DGAs for their C&C call-back communications.

New DGAs

Pleiades began clustering NXDomain traffic on November 1, 2010. We bootstrapped the DGA modeler with domain names from already known DGAs and also a set of Alexa domain names as the benign class. In Table 1, we present all unique clusters we discovered throughout the evaluation period. The "Malware Family" column simply maps the variant to a known malware family if possible. We discover the malware family by checking the NXDomains that overlap with NXDomains we extracted from traffic obtained from a malware repository. Also, we manually inspected the clusters with the help of a security company's threat

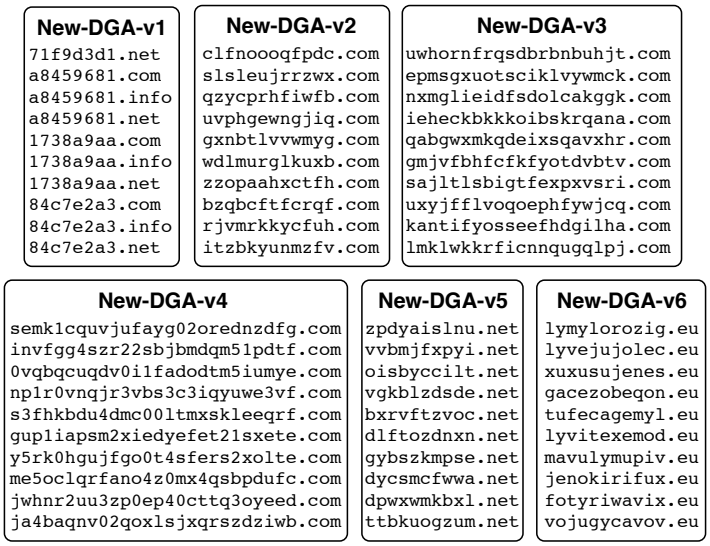

Figure 2: A sample of 10 NXDomains for each DGA cluster that we could not associate with a known malware family

Malware Family	First Seen	Population on Discovery
Shiz/Simda-C [8]	03/20/11	37
Bamital [4]	04/01/11	175
BankPatch [2]	04/01/11	28
Expiro.Z [3]	04/30/11	7
Boonana [9]	08/03/11	24
Zeus.v3 [7]	09/15/11	39
TDSS/TDL DGA Variant	07/08/12	201
New-DGA-v1	01/11/10	12
New-DGA-v2	01/18/11	10
New-DGA-v3	02/01/11	18
New-DGA-v4	03/05/11	22
New-DGA-v5	04/21/11	5
New-DGA-v6	11/20/11	10

Table 1: DGAs detected by Pleiades

team. The "First Seen" column denotes the first time we saw traffic from each DGA variant. Finally, the "Population on Discovery" column shows the variant population on the discovery day. We can see that we can detect each DGA variant with an average number of 32 "infected hosts" across the entire statewide ISP network coverage.

As we see in Table 1, Pleiades reported seven variants that belong to known DGA-enabled malware families [2–4, 7–9]. Six more variants of NXDomains were reported and modeled by Pleiades, but for these, to the best of our knowledge, no known malware can be associated with them. A sample set of 10 NXDomains for each one of these variants can be seen in Figure 2.

Within a two-year period of our experiments, we observed an average population of 742 Conficker-infected hosts in the ISP network. Murofet had the second largest population of infected hosts at 92 per day, while the Boonana DGA came in third with an average population of 84 infected hosts per day. The fastest growing DGA was Zeus.v3 with an average population of 50 hosts per day, but during the last four days of the experiments, the Zeus.v3 DGA had an average of 134 infected hosts. It is worth noting the New-DGA-v1 had an average of 19 hosts per day, the most populous of the newly identified DGAs.

FALSE REPORTS ON NEW DGAS

During our evaluation period we came across five categories of clusters falsely reported as new DGAs. In all of the cases, we modeled these classes in the DGA modeler as variants of the benign class. We now discuss each case in detail.

The first cluster of NXDomains falsely reported by Pleiades were random domain names generated by Chrome [10, 5]. Each time the Google Chrome browser starts, it will query three "random looking" domain names. These domain names are issued as a DNS check, so the browser can determine whether NXDomain rewriting is enabled. The "Chrome DGA" was reported as a variant of Bobax from Pleiades. We trained a class for this DGA and flagged it as benign. One more case

of testing for NXDomain rewriting was identified in a brand of wireless access points: Connectify offers wireless hot-spot functionality, and one of their configuration options enables the user to hijack the ISP's default NXDomain rewriting service. The device generates a fixed number of NXDomains to test for rewriting.

Two additional cases of false reports were triggered by domain names from the `.it` and `.edu` TLDs. These domain names contained minor variations on common words (i.e., repubblica, gazzetta, computer, etc.). Domain names that matched these clusters appeared only for two days in our traces and never again. The very short-lived presence of these two clusters could be explained if the domain names were part of a spam campaign that was remediated by authorities before it became live.

The fifth case of false report originated from domain names under a US government zone and contained the string `wpdhsmp`. Our best guess is that these are internal domain names that were accidentally leaked to the recursive DNS server of our ISP. Domain names from this cluster appeared only for one day. This class of NXDomains was also modeled as a benign variant. It is worth noting that all falsely reported DGA clusters, excluding the Chrome cluster, were short-lived. If operators are willing to wait a few days until a new DGA cluster is reported by Pleiades, these false alarms would not have been raised.

Case Studies

Next we discuss the three most interesting active threats that employ DGA techniques as part of their C&C life cycle.

ZEUS.V3

In September 2011, Pleiades detected a new DGA that we linked to the Zeus. v3 variant a few weeks later. The domain names collected from the machines compromised by this DGA-malware are hosted in six different TLDs: `.biz`, `.com`, `.info` , `.net` , `.org`, and `.ru`. Excluding the top-level domains, the length of the domain names generated by this DGA are between 33 and 45 alphanumeric characters. By analyzing one sample of the malware, we observed that its primary C&C infrastructure is P2P-based. If the malware fails to reach its P2P C&C network, it follows a contingency plan, where a DGA-based component is used to try to recover from the loss of C&C communication. The malware will then resolve pseudo-random domain names, until an active C&C domain name is found.

To date, we have discovered 12 such C&C domains. Over time, these 12 domains resolved to five different C&C IPs hosted in four different networks: three in the US (AS6245, AS16626, and AS3595) and one in the United Kingdom (AS24931). Interestingly, we observed that the UK-based C&C IP address remained active for only a few minutes, from Jan 25, 2012 12:14:04 EST to Jan 25, 2012 12:22:37 EST. The C&C moved from a US IP (AS16626) to the UK (AS24931), and then almost immediately back to the US (AS3595).

BANKPATCH

We picked the BankPatch DGA cluster as a sample case for analysis because this botnet had been active for several months during our experiments and the infected population continues to be significant. The C&C infrastructure that supports this botnet is impressive. Twenty-six different clusters of servers acted as the C&Cs for

this botnet. The botnet operators not only made use of a DGA but also moved the active C&Cs to different networks every few weeks (on average). During our C&C discovery process, we observed IP addresses controlled by a European CERT. This CERT has been taking over domain names from this botnet for several months. We managed to cross-validate with them the completeness and correctness of the C&C infrastructure. Complete information about the C&C infrastructure can be found in Table 2.

IP Addresses	CC	Owner
146.185.250.{89-92}	RU	Petersburg Int.
31.11.43.{25-26}	RO	SC EQUILIBRIUM
31.11.43.{191-194}	RO	SC EQUILIBRIUM
46.16.240.{11-15}	UA	iNet Colocation
62.122.73.{11-14,18}	UA	"Leksim" Ltd.
87.229.126.{11-16}	HU	Webenlet Kft.
94.63.240.{11-14}	RO	Com Frecatei
94.199.51.{25-18}	HU	NET23-AS 23VNET
94.61.247.{188-193}	RO	Vatra Luminoasa
88.80.13.{111-116}	SE	PRQ-AS PeRiQuito
109.163.226.{3-5}	RO	VOXILITY-AS
94.63.149.{105-106}	RO	SC CORAL IT
94.63.149.{171-175}	RO	SC CORAL IT
176.53.17.{211-212}	TR	Radore Hosting
176.53.17.{51-56}	TR	Radore Hosting
31.210.125.{5-8}	TR	Radore Hosting
31.131.4.{117-123}	UA	LEVEL7-AS IM
91.228.111.{26-29}	UA	LEVEL7-AS IM
94.177.51.{24-25}	UA	LEVEL7-AS IM
95.64.55.{15-16}	RO	NETSERV-AS
95.64.61.{51-54}	RO	NETSERV-AS
194.11.16.133	RU	PIN-AS Petersburg
46.161.10.{34-37}	RU	PIN-AS Petersburg
46.161.29.102	RU	PIN-AS Petersburg
95.215.{0-1}.29	RU	PIN-AS Petersburg
95.215.0.{91-94}	RU	PIN-AS Petersburg
124.109.3.{3-6}	TH	SERVENET-AS-TH-AP
213.163.91.{43-46}	NL	INTERACTIVE3D-AS
200.63.41.{25-28}	PA	Panamaserver.com

Table 2: C&C infrastructure for BankPatch

CIDR	CC	Owner
146.185.250.0/24	RU	PIN-AS
83.133.0.0/16	EU	LAMBDANET-AS
195.3.144.0/22	LV	RN-DATA-LV
94.63.149.0/24	RO	CORAL-IT
194.11.16.0/24	RU	PIN-AS
94.63.240.0/24	RO	POSTOLACHE
188.95.48.0/21	NL	GLOBALLAYER
46.251.224.0/20	DE	WEBTRAFFIC
95.215.0.0/22	RU	PIN-AS
94.60.122.0/23	RO	COVER-SUN-DESIGN
109.236.80.0/20	NL	WORLDSTREAM
63.223.96.0/19	US	JOVITA
91.212.226.0/24	RU	ZHIRK
46.161.28.0/23	RU	PIN-AS
141.136.16.0/24	RO	SC-MORE-SECURE-SRL
46.249.32.0/19	NL	SERVERIUS-AS
217.23.0.0/20	NL	WORLDSTREAM
62.122.74.0/23	EU	ROOT SA
50.7.192.0/19	US	FDCSERVERS
38.0.0.0/8	US	COGENT Cogent/PSI
194.247.182.0/23	UA	UDNET
195.234.124.0/22	UA	KOSMOTEL
195.28.10.0/23	RU	Neryungrinskoye
89.208.144.0/20	RU	DINET-AS
94.228.208.0/20	NL	NETROUTING-AS
27.255.64.0/19	KR	LGDACOM
91.199.75.0/24	DE	INTEROUTE
120.197.80.0/20	CN	CMNET

Table 3: Extended criminal network infrastructure behind New TDSS/TDL4 DGA variant

The actual structure of the domain name used by this DGA can be separated into a four-byte prefix and a suffix string argument. The suffix string arguments we observed were: `seapollo.com`, `tomvader.com`, `aulmala.com`, `apontis.com`, `fnomosk.com`, `erhogeld.com`, `erobots.com`, `ndsontex.com`, `rtehedel.com`, `nconnect.com`, `edsafe.com`, `berhogeld.com`, `musallied.com`, `newnacion.com`, `susaname.com`, `tvolveras.com`, `dminmont.com`, `esroater.com`, `jierihon.com` and `mobama.com`.

The four bytes of entropy for the DGA were provided by the prefix. We observed collisions between NXDomains from different days, especially when only one suffix argument was active. Therefore, we registered a small sample of 10 domain names at the beginning of 2012 in an effort to obtain a glimpse of the overall distribution of this botnet. Over a period of one month of monitoring the sinkholed data from the domain name of this DGA, this botnet has infected hosts in 270 different networks distributed across 25 different countries. By observing the recursive DNS servers from the domain names we sinkholed, we determined 4,295 were located in the US. The recursives we monitored were part of this list, and we were able to measure 86 infected hosts (on average) in the network we were monitoring. The five countries that had the most DNS resolution requests for the sinkholed domain names (besides the US) were Japan, Canada, the United Kingdom, and Singapore. The average number of recursive DNS servers from th ese countries that contacted our authorities was 22, significantly smaller than the volume of recursive DNS servers within the US.

TDSS/TDL4 DGA VARIANT

This TDSS/TDL4 DGA variant is the latest DGA discovery made possible by Pleiades. At the time of this writing, no malware sample has been discovered for this DGA variant. We believe that the DGA is primarily used to serve traditional C&C and enables click-fraud activities for the main TDSS/TDL4 [6] infection. This new DGA variant for TDSS/TDL4 appeared as a new DGA cluster in the beginning of July 2012. The C&C network-hosting infrastructure spans multiple different networks in Europe, US, and Asia. While most of the C&C IP addresses have been associated in the past with illicit operations (i.e., RBN, BitCoin mining) and have affected hundreds of thousands of victims, we are not aware of a sample available to the security community that resembles the DGA's behavior.

In an effort to describe the extended criminal C&C network for this TDSS/TDL4 variant, we first obtained the C&C domain names and remote IPs from the successful DNS resolutions observed by the TDSS/TDL4 DGA victims. Then we projected them in our passive DNS data collection in order to discover their immediate related historic IPs and domain names. We then selected all the domain names that matched the HMM model for this DGA variant. The resulting set of resource records constitutes the extended TDSS/TDL4 C&C network. Using the RDATA extracted from our passive DNS, we can provide a complete picture of the extended C&C network components. In Table 3, we show the extended criminal network behind this threat. We were able to identify 85 hosts that appear to be related to the actors behind TDSS/TDL4 DGA and that were used over the past 18 months.

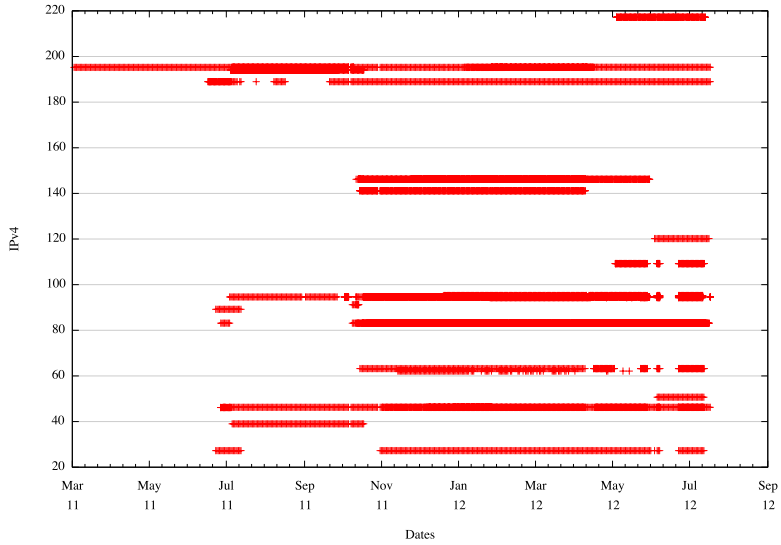

Figure 3: TDSS/TDL4 network agility

In Figure 3, we observe the network agility of the extended TDSS/TDL4 C&C network infrastructure and note how the botmasters behind TDSS/TDL4 moved and updated their impressive C&C network infrastructure from 03/03/2011 through 07/18/2012. Multiple C&Cs hosts were clearly active at the same time, especially towards the last few months of our analysis period.

In Figure 4, we present a small sample from the NXDomains the TDSS/TDL4 DGA generated over time. A few new NXDomains appear to be generated by the infected hosts every 48 hours. Using this observation, and in collaboration with Georgia Tech Information Security Center (GTISC), we managed to get a glimpse of the botnet worldwide infection levels. As of September 15, 2012, we have observed more than 250,000 unique Internet hosts tying to contact the GTISC sinkhole. Unfortunately, this number is growing, which implies that either the infection campaign is still active or the threat is largely undetectable by traditional network and host level defenses.

0sso151a47nztrxld6.com	ntibwlwhgl6bjp.com
0ubpccgkvzrnng.com	nxadrmwfgplgpr72jcv.com
1jndmf93bnobmbm-v.com	qiqxqoedngojyw4d.com
3tvcqyg4msj9byffzm.com	sbvv2b59psvpghaojz9.com
4rtgobtumvrzoqwq.com	smaug.gtisc.gatech.edu
ad9btvkonim6wsgg8lv.com	t407bqgh56jbkv4ua.com
anz7sjg6awufloz.com	udf-szhubujmuhp1jj.com
cudkkm05bzvn0dth.com	v-qk5nvogztncpmg2cp.com
d8kkkblaj6c4olp.com	vlxbbxrhq1llnft.com
dklfebjexiabttkwvgos.com	vpmybkeogu4vfuiu5s.com
fjg56xwoupqpdxr.com	wxbppgbdwiedzbnzhh.com
fwjudokrkhlwd3sm.com	xrqc-swswrykw30p.com
hrai41zpyw73sxhja5k3.com	yhftaw6wxlrhbl90osg.com
ikh9w-3vdmlndafja.com	ymgn1thqfe4q6rs.com
nihawelnopjmpn67yrn.com	zv7dfcgtusnttpl.com

Figure 4: TDSS/TDL4 DGA NXDomain samples

Conclusion

With this short article, we summarize the key aspects of a novel detection system called Pleiades. This system is able to detect machines accurately within a monitored network that are compromised with DGA-based bots. Utilizing the streams of unsuccessful DNS resolution from a large ISP, Pleiades can identify and model previously unknown DGAs, instead of relying on manual reverse engineering of bot malware and their DGA algorithms. In our multi-month evaluation phase, Pleiades was able to identify seven DGAs that belong to known malware families and six new DGAs never reported before.

References

[1] M. Antonakakis, R. Perdisci, Y. Nadji, N. Vasiloglou, S. Abu-Nimeh, W. Lee, and D. Dagon, "From Throw-Away Traffic to Bots: Detecting the Rise of DGA-Based Malware," 21st USENIX Security Symposium, USENIX Association, 2012.

[2] BankPatch, Trojan.Bankpatch.C: http://www.symantec.com/security_response/writeup.jsp?docid=2008-081817-1808-99, 2009.

[3] Microsoft Malware Protection Center, Virus:Win32/Expiro.Z: http://www.microsoft.com/security/portal/Threat/Encyclopedia/Entry.aspx?Name=Virus%3AWin32%2FExpiro.Z, 2011.

[4] M. Geide, "Another Trojan Bamital Pattern": http://research.zscaler.com/2011/05/another-trojan-bamital-pattern.html, 2011.

[5] S. Krishnan and F. Monrose, "DNS Prefetching and Its Privacy Implications: When Good Things Go Bad," *Proceedings of the 3rd USENIX Conference on Large-Scale Exploits and Emergent Threats* (LEET '10), USENIX Association, 2010, pp. 10-10.USENIX Association.

[6] A. Matrosov, E. Rodionov, and D. Harley, TDSS parts 1 through 3: http://resources.infosecinstitute.com/tdss4-part-1/, http://resources.infosecinstitute.com/tdss4-part-2/, http://resources.infosecinstitute.com/tdss4-part-3/, 2011.

[7] CERT Polska, "ZeuS P2P+DGA Variant Mapping Out and Understanding the Threat": http://www.cert.pl/news/4711/langswitch_lang/en, 2012.

[8]Sophos, Mal/Simda-C: http://www.sophos.com/en-us/threat-center/threat-analyses/viruses-and-spyware/Mal~Simda-C/detailed-analysis.aspx, 2012.

[9] Microsoft Malware Protection Center, Trojan:Java/Boonana: http://www.microsoft.com/security/portal/Threat/Encyclopedia/Entry.aspx?Name=Trojan%3AJava%2FBoonana, 2011.

[10] B. Zdrnja, "Google Chrome and (Weird) DNS Requests": http://isc.sans.edu/diary/Google+Chrome+and+weird+DNS+requests/10312, 2011.