

Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Honey pot trace forensics: The observation viewpoint matters

Van-Hau Pham^{a,*}, Marc Dacier^b

^a School of Computer Science & Engineering, International University, Hochiminh City, Viet Nam

^b Symantec Research Labs, Sophia Antipolis, France

ARTICLE INFO

Article history:

Received 30 November 2009

Received in revised form

29 April 2010

Accepted 14 June 2010

Available online 23 June 2010

Keywords:

Honey pot

Attack trace analysis

Botnet detection

ABSTRACT

In this paper, we propose a method to identify and group together traces left on low interaction honeypots by machines belonging to the same botnet(s) without having any a priori information at our disposal regarding these botnets. In other words, we offer a solution to detect new botnets thanks to very cheap and easily deployable solutions. The approach is validated thanks to several months of data collected with the worldwide distributed Leurré.com system. To distinguish the relevant traces from the other ones, we group them according to either the platforms, i.e. targets hit or the countries of origin of the attackers. We show that the choice of one of these two observation viewpoints dramatically influences the results obtained. Each one reveals unique botnets. We explain why. Last but not the least, we show that these botnets remain active during very long periods of times, up to 700 days, even if the traces they left are only visible from time to time.¹

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

There is a consensus in the security community to say that botnets are today's plague of the Internet. A lot of attention has been paid to detect and eradicate them. Several approaches have been proposed for this purpose. By identifying the so-called *Command and Control* (C&C) channels, one can keep track of all IPs connecting to it. The task is more or less complicated, depending on the type of C&C (IRC [1–4], HTTP [5,6], fast-flux based or not [7,8], P2P [9–11], etc.) but, in any case, one needs to have some insight about the channels and the capability to observe all communications on them. Another approach consists in sniffing packets on a network and in recognizing patterns of *bot-like* traffic. This is, for instance, the approach pursued by [12–15]. The solutions mostly aim at detecting compromised machines in a given network rather than to study the botnets themselves as they only see the bots that exist within the network under study.

In this work, we are interested in finding a very general technique that would enable us to count the amount of various botnets that exist, their size and their lifetime. As opposed to previous work, we are not interested in studying a particular botnet in detail or in detecting compromised nodes in a given network. We also do not want to learn the various protocols used by bots to communicate in order to infiltrate the botnets and

obtain more precise information about them [4]. By doing so, we certainly will not be able to get as much in depth information about this or that botnet but our hope is to provide insights into the bigger picture of today's (and yesterday's) botnet activities. This kind of knowledge could be used by defenders when designing the countermeasures.

The solution described in the following is generic and simple to deploy widely. It relies on a distributed system of low interaction honeypots. Based on the traces left on these honeypots, we provide a technique that groups together the traces that are likely to have been generated by groups of machines controlled by a similar authority. Since we have no information regarding the C&C they obey, we do not know if these machines are part of a single botnet or if they belong to several botnets that are coordinated. Therefore, to avoid any ambiguity, we write in the following that they are part of an *army of zombies*. An *army of zombies* can be a single botnet or a group of botnets the actions of which are coordinated during a given time interval.

In this paper, we propose a technique to identify and study the size as well as the lifetime of such *armies of zombies*. The approach does not pretend to be able to identify all *armies of zombies* that could be found in our dataset. On the contrary, we show that, depending on how the dataset is preprocessed, i.e. depending on the observation viewpoint, different armies can be found. Exhaustiveness is not our concern at this stage but, instead, we are interested in offering an approach that could easily be widely adopted.

The idea exposed here is similar, in its spirit, to the one presented in the paper coauthored by Allman et al. [16]. However, instead of “[...] leveraging the deep understanding of network

* Corresponding author.

E-mail addresses: pvhau@hcmiu.edu.vn, vanhau.pham@gmail.com (V.-H. Pham), Marc_Dacier@symantec.com (M. Dacier).

¹ The present paper is an extended version of Pham and Dacier (2009) [25].

detectives and the broad understanding of a large number of network witnesses to form a richer understanding of large-scale coordinated attackers”, our approach relies on a diverse yet limited number of low interaction honeypots. They do not need to be neither as smart as the network detectives nor as numerous as the network witnesses proposed in that work. Both approaches are quite complementary. Kitti et al. have proposed an approach to detect *related attacks* in [17]. The method has been validated thanks to data collected from DShield project [18]. In that work, related attacks are understood as attacks mounted by the same sources against different networks which is a narrower view of the problem than ours.

Finally, our approach is also different from the one adopted in [19]. In fact, in [19], the botnet detection module must be installed within the networks where bots reside to detect them whereas, in our case, our honeypots are the targets of the attacks.

The remainder of the paper is organised as follows. Section 2 defines the terms used in the paper. Section 3 describes the dataset we have used and what we mean when we refer to the notion of *observation viewpoint*. It provides some motivation for the work. In Section 4, we describe the method itself and provide the main characteristics of the results obtained as well as two precise, yet anecdotal, examples of armies detected thanks to our method. Finally, Section 5 concludes the paper.

2. Terminology

In order to avoid any ambiguity, we introduce a few terms that will be used throughout the text. Some of them are taken from [20]. Readers who are familiar with the Leurré.com project are invited to skip this Section.

- **Platform:** A physical machine simulating, thanks to honeyd [21], the presence of three distinct machines. A platform is connected directly to the Internet and collects tcpdump traces that are fed daily into the centralized Leurré.com’s database.
- **Leurré.com:** The Leurré.com project is a distributed system of such platforms deployed in more than 50 different locations in 30 different countries (see [22] for details).
- A **Source** corresponds to an IP address that has sent at least one packet to, at least, one platform. A given IP address can correspond to several distinct sources. Indeed, a given IP remains associated to a given source as long as there is no more than 25 h between 2 packets received from that IP. After that, a new source identifier will be assigned to the IP. By grouping packets by sources instead of by IPs, we minimize the risk of gathering packets sent by distinct physical machines that have been assigned the same IP dynamically after 25 h, or machines that have the same IP address seen from the outside due to side effect of Network Address Translation.
- An **Attack**, in the context of this paper, is defined as the packets exchanged between one source and one platform.
- A **Cluster** is made of a group of sources that have left highly similar network traces on all platforms they have been seen on. Clusters have been precisely defined in [23].
- An **Observed cluster time series** $\Phi_{T,c,op}$ is a function defined over a period of time T , T being defined as a time interval (in days). That function returns the amount of sources per day associated to a cluster c that can be seen from a given *observation viewpoint* op . The observation viewpoint can either be a specific platform or a specific country of origin. In the first case, $\Phi_{T,c,platform_x}$ returns, per day, the amount of sources belonging to cluster c that have hit $platform_x$. Similarly, in the second case, $\Phi_{T,c,country_x}$ returns, per day, the amount of sources belonging to cluster c that are geographically located in $country_x$. Clearly, we always have: $\Phi_{T,c} = \sum_{\forall i \in countries} \Phi_{T,c,i} = \sum_{\forall x \in platforms} \Phi_{T,c,x}$.

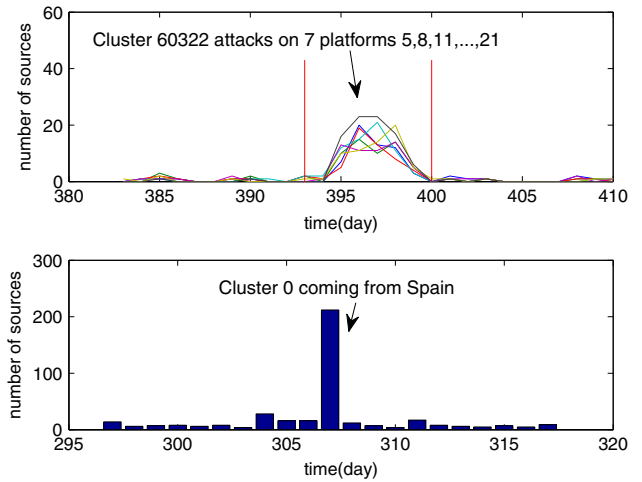


Fig. 1. On the top plot, cluster 60,232 attacks seven platforms from day 393 to day 400. On the bottom plot, peak of activities of cluster 0 from Spain on day 307.

- An **attack event** is defined as a set of observed cluster time series exhibiting a particular shape during a limited time interval. The set can be a singleton. We denote the attack event i as $e_i = (T_{start}, T_{end}, S_i)$ where the attack event starts at T_{start} , ends at T_{end} and S_i contains a set of observed cluster time series identifiers (c_i, op_i) such that all $\Phi_{[T_{start}-T_{end}],c_i,op_i}$ are strongly correlated to each other $\forall (c_i, op_i) \in S_i$. As an example, the top plot of Fig. 1 represents the attack event 225 which consists of a given cluster attacking seven platforms. Each curve represents the amount of sources of that cluster observed from one of these platforms. As we can observe, the attack event starts at day 393 and ends at day 400. According to our convention, we have $e_{225} = (393, 400, \{(60232, 5), (60232, 8), \dots, (60232, 31)\})$. Similarly, the bottom plot of Fig. 1 represents an attack event due to one cluster during a single day and mostly due to a single country ($e_{14} = (307, 307, \{(0, ES)\})$).

3. Impact of the observation viewpoint

3.1. Dataset description

For our experiments, we have selected the traces observed on 40 platforms out of 50 at our disposal. All these 40 platforms have been running for more than 800 days. None of them has been down for more than 10 times and each of them has been up continuously for at least 100 days at least once. They all have been up for a minimum of 400 days over that period. We denote by T , the time series representing the total amount of sources observed, day by day, on all these 40 platforms. We can split that time series per country² of origin of the sources. This gives us 231 time series TS_x where the i th point of such time series indicates the amount of sources, observed on all platforms, located in country X . We represent by TS_{L1} the set of all these Level 1 time series. To reduce the computational cost, we keep only the countries from which we have seen at least 10 sources on at least one day. This leaves us with 85, instead of 231, time series. We represent by $TS_{L1'}$ this refined set of Level 1 time series. Then, we split each of these time series by cluster to produce the final set of time series $\Phi_{[0-800],c_i,country_j} \forall c_i$ and $\forall country_j \in big_{countries}$. The i th point of the time series $\Phi_{[0-800],X,Y}$ indicates the amount of sources originating from country Y that has been observed on day i attacking any of our

² We use Maxmind to get the geographical location of IPs.

Table 1

Dataset description: TS : all sources observed on the period under study, OVP : observation viewpoint, TS_L1 : set of time series at country/platform level, TS_L1' : set of significant time series in TS_L1 , TS_L2 : set of all cluster time series, TS_L2' : set of strongly varying cluster time series.

TS consists of 3477,976 sources		
OVP	Country	Platform
$ TS_L1 $	231	40
$ TS_L1' $	85	40
	(944% TS)	(100% TS)
$ TS_L2 $	436,756	395,712
$ TS_L2' $	2420	2127
sources	2330,244	2538,922
	(67% of TS)	(73% of TS)

platforms thanks to the attack defined by means of the cluster X . We represent by TS_L2 the set of all these Level 2 time series. In this case $|TS_L2|$ is equal to 436,756 which corresponds to 3284,551 sources.

As explained in [20], time series that barely vary in amplitude over the 800 days are meaningless to identify attack events and we can get rid of them. Therefore, we only keep the time series that highlight important variations. We represent by TS_L2' this refined set of Level 2 time series. In this case $|TS_L2'|$ is equal to 2420 which corresponds to 2,330,244 sources.

We have done the very same splitting and filtering by looking at the traces on a per platform basis instead of on a per country of origin basis. The corresponding results are given in Table 1.

3.2. Attack event detection

Having defined the time series we are interested in, we now need to identify all time periods during which 2 or more of these observed cluster time series are correlated together.

To do this, in a first step, we use a sliding window of L days to compute the Pearson correlation of all pairs of time series. That is, we compute the correlation of N time series for $T - L + 1$ time interval $\{[1, L], [2, L + 1], \dots, [T - L, T]\}$. As a result, we obtain, for every pair of time series in N , the time intervals during which they are correlated. Then we group together all pairs of cluster time series that are correlated together over the same period of time. Each such group constitutes an *attack event* as defined before.

It is worth noting that this method, which we refer to as $M1$ in the sequel, cannot detect attack events made of a single cluster time series. This is typically the case for peaks of activities occurring on a single day. In such cases, it is more efficient to apply another, less expensive, algorithm to identify the attack events. For the sake of conciseness, we do not to include the description of this second method, $M2$.

3.3. Impact of the observation viewpoint

3.3.1. Results on attack event detection

We have applied these algorithms against our 2 distinct datasets, namely $TS_{country}$ and $TS_{platform}$. As shown in Table 2, for $TS_{country}$, method $M1$ (resp. second method $M2$) has found 549 (resp. 43) attack events, accounting for a total of 552,492 sources (resp. 21,633). Similarly, with $TS_{platform}$, applying $M1$ (resp. $M2$) leads to 564 (resp. 126) attack events, containing 550,305 (resp. 28,067) sources.

3.3.2. Analysis

The table highlights the fact that depending on how we decompose the initial set of traces of attacks (i.e the initial time

Table 2

Result on attack event detection.

	AE-set-I($TS_{country}$)		AE-set-II($TS_{platform}$)	
	No. AEs	No. sources	No. AEs	No. sources
M1	549	552,492	564	550,305
M2	43	21,633	126	28,067
Total	592	574,125	690	578,372

No.AEs: amount of attack events. M1, M2: methods represented in Section 3.2.




Fig. 2. CDF common source ratio.

series TS), namely by splitting it by countries of origin of the attackers or by platforms attacked, different attacks events show up. To assess the overlap between attack events detected from different observation viewpoints we use the *common source ratio*, namely csr , measure as follows:

$$csr(e, AE_{op'}) = \frac{\sum_{e' \in AE_{op'}} |e \cap e'|}{|e|}$$

in which $e \in AE_{op}$ and $|e|$ is the amount of sources in attack event e , AE_{op} is $AE_{country}$ and $AE_{op'}$ is $AE_{platforms}$ (or vice versa).

Fig. 2 represents the two cumulative distribution functions corresponding to this measure. The point (x, y) on the curve means that there are $y * 100\%$ of attack events obtained thanks to $T_{country}$ (resp $T_{platforms}$) that have less than $x * 100\%$ of sources in common with all attack events obtained thanks to $T_{platforms}$ (resp $T_{country}$). The $T_{country}$ curve represents the cumulative distribution obtained in this first case and the $T_{platforms}$ one represents the CDF obtained when starting from the attacks events obtained with the initial $T_{platforms}$ set of time series. As we can notice, around 23% (resp. 25%) of attack events obtained by starting from the $T_{country}$ (resp. $T_{platform}$) set of time series do not share any source in common with any attack events obtained when starting the attack even identification process from the $T_{platform}$ (resp. $T_{country}$) set of time series. This corresponds to 136 (16,919 sources) and 171 (75,920 sources) attack events not being detected. In total, there are 288,825 (resp. 293,132) sources present in AE-Set-I (resp. AE-Set-II), but not in AE-Set-II (resp. AE-Set-I). As a final note, there are in total 867,248 sources involved in all the attack events detected from both datasets which correspond to 25% the attacks observed in the period under study.

3.3.3. Explanation

The reasons why we cannot rely on a single viewpoint to detect all attacks events are described below.

Split by country: Suppose we have one botnet B made of machines that are located within the set of countries $\{X, Y, Z\}$. Suppose

that, from time to time, these machines attack our platforms leaving traces that are also assigned to a cluster C . Suppose also that this cluster C is a very *popular* one, that is, many other machines from all over the world continuously leave traces on our platforms that are assigned to this cluster. As a result, the activities specifically linked to the botnet B are lost in the noise of all other machines leaving traces belonging to C . This is certainly true for the cluster time series (as defined earlier) related to C and this can also be true for the time series obtained by splitting it by platform, $\Phi_{[0-800],C,platform_i} \forall platform_i \in 1 \dots 40$. However, by splitting the time series corresponding to cluster C by countries of origins of the sources, then it is quite likely that the time series $\Phi_{[0-800],C,country_i} \forall country_i \in \{X, Y, Z\}$ will be highly correlated during the periods in which the botnet present in these countries will be active against our platforms. This will lead to the identification of one or several attack events.

Split by platform: Similarly, suppose we have a botnet B' made of machines located all over the world. Suppose that, from time to time, these machines attack a specific set of platforms $\{X, Y, Z\}$ leaving traces that are assigned to a cluster C . Suppose also that this cluster C is a very *popular* one, that is, many other machines from all over the world continuously leave traces on all our platforms that are assigned to this cluster. As a result, the activities specifically linked to the botnet B' are lost in the noise of all other machines leaving traces belonging to C . This is certainly true for the cluster time series (as defined earlier) related to C and this can also be true for the time series obtained by splitting it by countries, $\Phi_{[0-800],C,country_i} \forall country_i \in big_{countries}$. However, by splitting the time series corresponding to cluster C by platforms attacked, then it is quite likely that the time series $\Phi_{[0-800],C,platform_i} \forall platform_i \in \{X, Y, Z\}$ will be highly correlated during the periods in which the botnet influences the traces left on the sole platforms concerned by its attack. This will lead to the identification of one or several attack events.

The top plot of Fig. 3 represents the attack event 79. In this case, we see that the traces due to the cluster 175,309 are highly correlated when we group them by platform attacked. In fact, there are 9 platforms involved in this case, accounting for a total of 870 sources. If we group the same set of traces by country of origin of the sources, we end up with the bottom curves of Fig. 3 where the specific attack event identified previously can barely be seen. This highlights the existence of a botnet made of machines located all over the world that target a specific subset of the Internet.

4. On the armies of Zombies

So far, we have identified what we have called attack events which highlight the existence of coordinated attacks launched by a group of compromised machines, i.e. a zombie army. It would be interesting to see if the very same army manifests itself in more than one attack event. To do this, we propose to compute what we call the *action sets*. An *action set* is a set of attack events that are likely due to the same army. In this Section, we show how to build these action sets and what information we can derive from them regarding the size and the lifetime of the zombie armies.

4.1. Identification of the armies

4.1.1. Similarity measures

In its simplest form, a zombie army is a classical botnet. It can also be made of several botnets, that is several groups of machines listening to distinct C 's. This is invisible to us and irrelevant. What matters is that all the machines do act in a coordinated way. As time passes, it is reasonable to expect members of an army to be cured while others join. So, if the same army attacks our honeypots twice




Fig. 3. The top plot represents the attack event 79 related to cluster 17,309 on 9 platforms. The bottom plot represents the evolution of this cluster by country. Noise of the attacks to other platforms decrease significantly the correlation of observed cluster time series when split by country.

over distinct periods of time, one simple way to link the two attack events together is by noticing that they have a large amount of IP addresses in common. More formally, we measure the likelihood of two attacks events e_1 and e_2 to be linked to the same army by means of their similarity defined as follows:

$$sim(e_1, e_2) = \begin{cases} \max\left(\frac{|e_1 \cap e_2|}{|e_1|}, \frac{|e_1 \cap e_2|}{|e_2|}\right) & \text{if } |e_1 \cap e_2| < 200 \\ 1 & \text{otherwise.} \end{cases}$$

We will say that e_1 and e_2 are caused by the same army if and only if $sim(e_1, e_2) > \delta$. This only makes sense for *reasonable* values of δ . We address this issue in the next subsections.

4.1.2. Action sets

We now use the $sim()$ function to group together attack events into action sets. To do so, we build a simple graph where the nodes are the attack events. There is an arc between two nodes e_1 and e_2 if and only if $sim(e_1, e_2) > \delta$. All nodes that are connected by at least one path end up in the same action set. In other words, we have as many action sets as we have disconnected graphs made of at least two nodes; singleton sets are not counted as action sets.

We note that our approach is such that we can have an action set made of three attack events e_1, e_2 and e_3 where $sim(e_1, e_2) > \delta$ and $sim(e_2, e_3) > \delta$ but where $sim(e_1, e_3) < \delta$. This is consistent with our intuition that armies can evolve over time in such a way that the machines present in the army can, eventually, be very different from the ones found the first time we have seen the same army in action.

4.1.3. Results

To test the sensitivity of the threshold δ , we have computed the amount of action sets for the two datasets for different values of δ . The result is represented in top plot of Fig. 4 (the bottom plot represent the corresponding amount of attack events involved in the armies). As we can see, at first, for the value of δ from 1% to 7%, the amount of action sets increases rapidly. Indeed, for very small values of δ all nodes remain connected together but, as δ increases, the initial graph loses arcs and more disconnected graphs appear, i.e. more action sets show up. This creation of action sets reaches a maximum after which action sets start disappearing with a growing δ value. This is due to the fact that some graphs are broken into isolated nodes that are not counting as attack sets




Fig. 4. Sensitivity check of threshold δ .

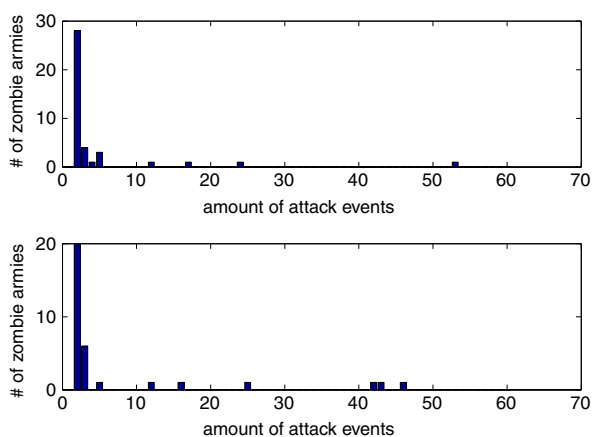


Fig. 5. Zombie army size.

anymore. The two curves reach their maximum values almost at the same position (when $\delta = 8\%$). Then they both start decreasing linearly.

A closer look shows that the threshold of $\delta = 10\%$ gives a good result to show in this paper. We do not pretend that this number is optimal in any sense and, in fact, we do not really care. Indeed, our purpose, at this stage, is just to look at the results for one given value of δ and see if, yes or no, this theory of zombie armies seems to be valid or not, based on the characteristics of the ones we will find in that particular case. It can very well be that the attack events found in attack sets, as we have built them, have no underlying common cause and that they accidentally share common IPs. In this paper, results presented have been obtained with a value of $\delta = 10\%$. Other values could, possibly, have delivered more armies but the point we want to make is that these armies exist, not that we have found a method to find all of them. For such value of δ we have identified 40 (resp. 33) zombie armies from AE-set-I (resp. AE-set-II) which have issued a total of 193 (resp. 247) attack events. Fig. 5 represents the distribution of attack events per zombie army. Its top (resp. bottom) plot represents the distribution obtained from AE-set-I (resp. AE-set-II). We can see that the largest amount of attack events for an army is 53 (resp. 47) whereas 28 (resp. 20) armies have been observed only two times.

4.2. Main characteristics of the zombie armies

In this section, we will analyze the main characteristic of the zombie armies.

Lifetime of Zombie Army. Fig. 6 represents the cumulative distribution of minimum lifetime of zombie armies obtained from $TS_{platform}$




Fig. 6. CDF duration.

and $TS_{country}$ (see Section 4.1.3). According to the plot, around 20% of zombie armies have existed for more than 200 days. In the extreme case, two armies seems to have survived for 700 days! Such result seems to indicate that either (i) it takes a long time to cure compromised machines or that (ii) armies are able to stay active for long periods of time, despite the fact that some of their members disappear, by continuously compromising new ones.

Lifetime of Infected Host in Zombie Armies. In fact, we can classify the armies into two classes as mentioned in the previous Section. For instance, Fig. 7a represents the similarity matrix of zombie army 33, ZA-33. To build this matrix, we first order its 42 attack events according to their time of occurrence. Then we represent their similarity relation under an 42×42 similarity matrix \mathcal{M} . The cell (i, j) represents the value of $sim()$ of the ordered attack event i th and j th. Since, \mathcal{M} is a symmetric matrix, only half of it is shown. As we can see, we have a very high similarity measure between almost all the attacks events, around 60%. This is also true between the very first and the very last attack events. In this case, the time elapsed between the first and the last event is 753 days!

Fig. 7(b) represents an opposite case, the zombie army 31, ZA31, consisting of 46 attack events. We proceed as above to build its similarity matrix. The important values are now located around the main diagonal of \mathcal{M} . It means that the attack event i th has the same subset of infected machines with only few attack events happening just before and after it. In this case, this army changed its attack vector over time, launching first attacks against 4662 TCP, then 1025 TCP, then 5900 TCP, 1443 TCP, 2967 TCP, 445 TCP, etc. Its lifetime is 563 days!

Attack Capacity. By attack capacity, we refer to the amount of different attacks that a given army is observed launching over time. The advanced worm, namely multi-headed worm, we have presented in our earlier work [20] is an example of worms that have many attack vectors and use them dynamically. The multi-attack vectors allow the worms to have a large chance to propagate, and the varying in activity helps them to have multi-attack traces which make it harder for IDS to detect them. This work reinforces the results we have earlier [20]. In fact, in previous work, we were able to detect multi-headed worms by the correlation of attack traces generated by different attack tools within an attack event. In this work, we have some even stronger evidence. Indeed, thanks to the notion of army, we observe several cases in which the same IP address has different behaviors in different attack events attached to a given army. As an example, the two attack events 128 and 131 consist of clusters 1378 and 2666 respectively. They both have 106 IP addresses in common and belong to the zombie army 12. All the attacks of attack event 128 are against port 64,783 TCP whereas all the attacks of attack event 131 are against port 6211 TCP. The conclusion is that these 106 attacking machines mentioned earlier have dynamically changed their behavior. Finally, Fig. 8 represents

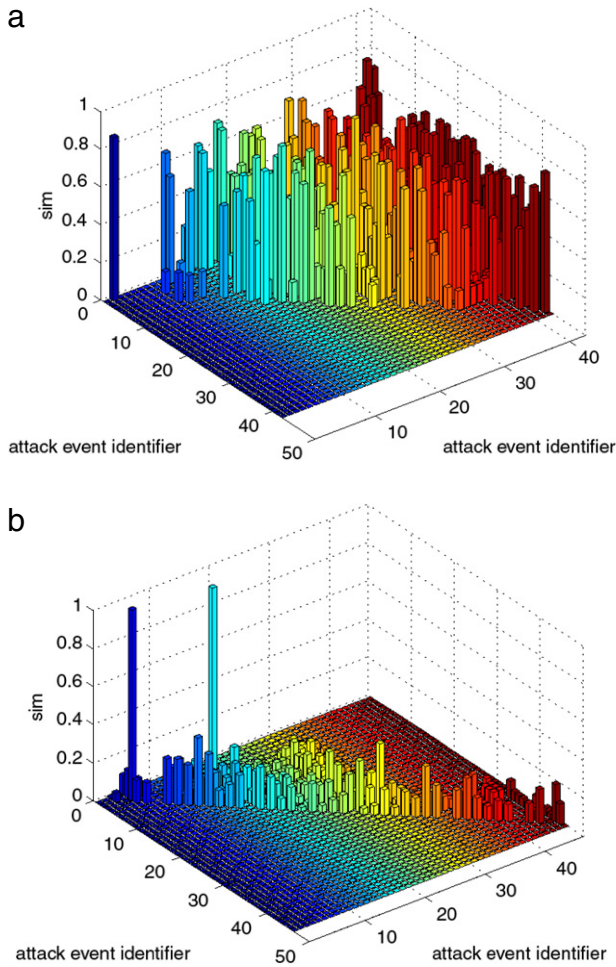


Fig. 7. Renewal rate of zombie armies.

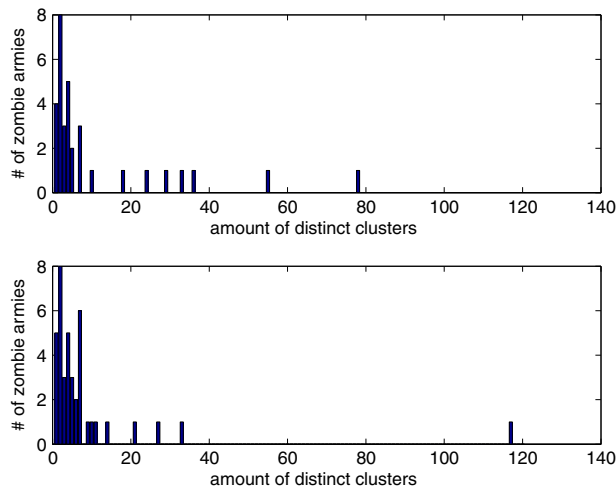


Fig. 8. Zombie army attack capacity.

the distribution of number of distinct cluster per army. One zombie army has almost 120 clusters. The large amount of distinct clusters can be due to the side effect of attack tools that have several attack scenarios, but it is more likely due to the update behavior of botnets. In fact, as observed in [24], “[...] The botmasters appear to ask most of the bots in a botnet to focus on one vulnerability, while choosing a small subset of the bots to test another vulnerability”.

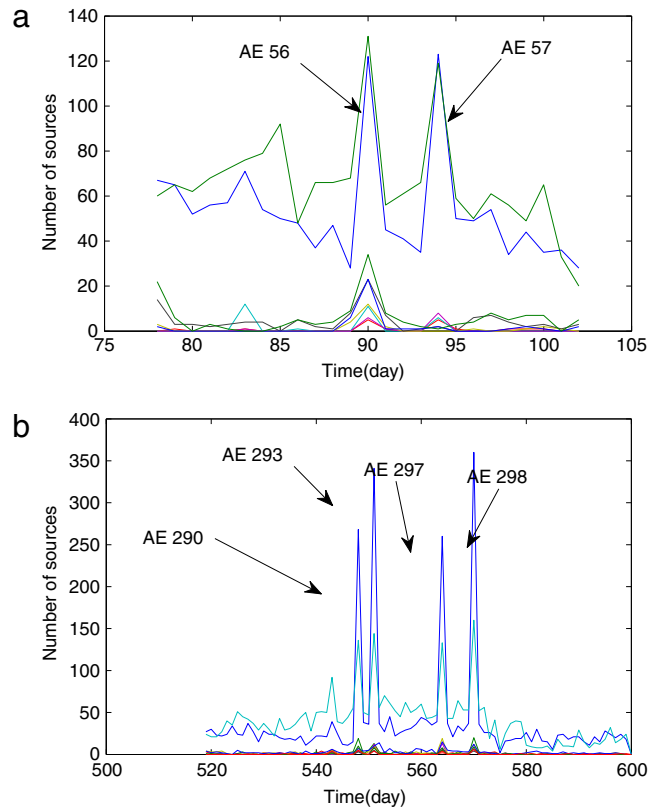


Fig. 9. Attack events of ZA-29.

4.3. Illustrated examples

After having offered a high level overview of the method and main characteristics of the results obtained, we feel it is important to give a couple of concrete, simple, examples of armies we have discovered. This should help the reader in better understanding the reality of two armies as well as what they look like. This is what we do in the next two subsections where we briefly present two representative armies.

4.3.1. Example 1

Zombie army 29, ZA-29, is an interesting example which has only been observed attacking a single platform. However, 16 distinct attack events are linked to that army! Fig. 9a presents its two first activities corresponding to the two attack events 56 and 57. Fig. 9b represents the other four attack events. In each attack event, the army tries a number of distinct clusters such as 13,882, 14,635, 14,647, 56,608, 144,028, 144,044, 149,357, 164,877, 166,477. These clusters try many combinations of Windows ports (135 TCP, 139 TCP, 445 TCP) and Web server (80 TCP). The time interval between the first and the last activities is 616 days !

4.3.2. Example 2

The zombie army 33, ZA-33, consisting of 42 attack events (already mentioned in Section 4.2) is an example of a multi-botnets zombies army. In fact, it seems that several botnets do different jobs and from time to time, they do some tasks together. In fact, in some cases, an important fraction of the machines in the attack events come from Italy and attack a single platform located in China. The two top plots in Fig. 10 represent such cases. The attack event 291 consists of several clusters attacking port 64,783 TCP. The attack event 195 also is mostly made of Italian sources and also uniquely target a platform in China but it is made of

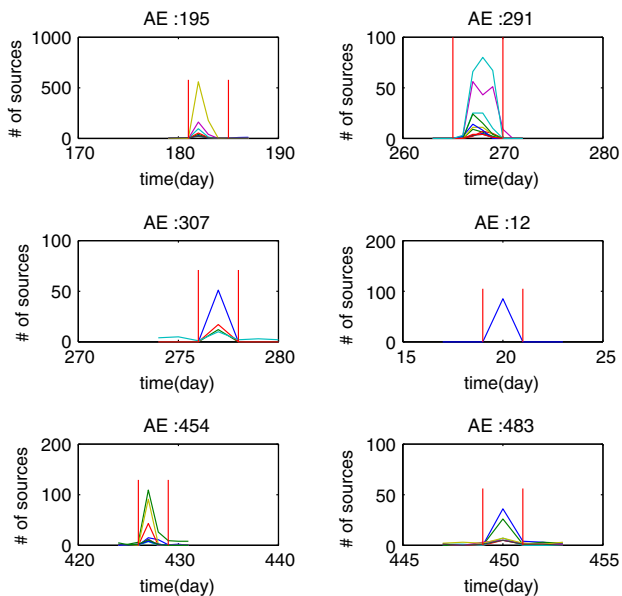


Fig. 10. 6 attack events from zombie army 33.

several clusters targeting port 9661 TCP. Interestingly enough, in some other cases, other attack events of the same army ZA-33 consistently sends ICMP packets only, are made of Greek sources, targeting a single platform also located in Greece (see the two plots in the middle of Fig. 10). As an example of coordination of two components of ZA-33, the two plots in the bottom of Fig. 10 represent two attack events (out of four) coming mostly from these two countries and attacking these two platforms. As a reminder, by design, there always is an overlap in terms of IP sources between the attack events. For instance, attack event 483 has 41 IP addresses in common with AE 307, whereas 454 and 483 have 47 IP addresses in common. The interval between the first and the last attack event issued by this zombie army is 753 days.

5. Conclusion

In this paper, we have addressed the important attack attribution problem. We have shown how low interaction honeypots can be used to track armies of zombies and characterize their lifetime and size. More precisely, this paper offers three main contributions. First of all, we propose a simple technique to identify, in a systematic and automated way, the so-called attack events in a very large dataset of traces. We have implemented and demonstrated experimentally the usefulness of this technique. Secondly, we have shown how, by grouping these attack events, we can identify long living armies of zombies. Here too, we have validated experimentally the soundness of the idea as well as the meaningfulness of the results it produces. Last but not least, we have shown the importance of the selection of the observation viewpoint when trying to group such traces for analysis purposes. Two such viewpoints have been considered in this paper, namely the geolocation of the attackers and the platform attacked. Results of the experiments have highlighted the benefits of considering more than one viewpoint as each of them offers unique insights into the attack processes. Future work includes the application of these techniques to richer data feeds, such as the ones produced by the European WOMBAT project (www.wombat-project.eu).

Acknowledgement

This work has been partially supported by the European Commissions through project FP7-ICT-216026-WOMBAT funded

by the 7th framework program. The opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the European Commission.

References

- [1] E. Cooke, F. Jahanian, D. McPherson, The zombie roundup: understanding, detecting, and disrupting botnets, in: SRUTT'05: Proceedings of the Steps to Reducing Unwanted Traffic on the Internet on Steps to Reducing Unwanted Traffic on the Internet Workshop, Berkeley, CA, USA: USENIX Association, 2005.
- [2] P. Barford, V. Yegneswaran, An inside look at botnets, in: Advances in Information Security, vol. 27, 2007, pp. 171–191.
- [3] J. Goebel, T. Holz, Rishi: identify bot contaminated hosts by irc nickname evaluation, in: Workshop on Hot Topics in Understanding Botnets 2007, 2007.
- [4] M. Rajab, J. Zarfoss, F. Monrose, A. Terzis, A multifaceted approach to understanding the botnet phenomenon, in: ACM SIGCOMM/USENIX Internet Measurement Conference, October 2006.
- [5] K. Chiang, L. Lloyd, A case study of the rustock rootkit and spam bot, in: First Workshop on Hot Topics in Understanding Botnets, 2007.
- [6] N. Daswani, M. Stoppelman, The anatomy of clickbota, in: HotBots'07: Proceedings of the First Workshop on Hot Topics in Understanding Botnets, Berkeley, CA, USA: USENIX Association, 2007.
- [7] T. Holz, C. Gorecki, K. Rieck, F.C. Freiling, Measuring and detecting fast-flux service networks, in: NDSS 2008, 2008.
- [8] E. Passerini, R. Paleari, L. Martignoni, D. Bruschi, Fluxor: detecting and monitoring fast-flux service networks, in: DIMVA 2008, 2008.
- [9] T. Holz, M. Steiner, F. Dahl, E. Biersack, F. Freiling, Measurements and mitigation of peer-to-peer-based botnets: a case study on storm worm, in: LEET'08: Proceedings of the 1st Usenix Workshop on Large-Scale Exploits and Emergent Threats, Berkeley, CA, USA: USENIX Association, 2008, pp. 1–9.
- [10] J.B. Grizzard, V. Sharma, C. Nunnery, B.B. Kang, D. Dagon, Peer-to-peer botnets: overview and case study, in: HotBots'07: Proceedings of the first conference on First Workshop on Hot Topics in Understanding Botnets, Berkeley, CA, USA: USENIX Association, 2007.
- [11] P. Wang, S. Sparks, C.C. Zou, An advanced hybrid peer-to-peer botnet, in: HotBots'07: Proceedings of the first conference on First Workshop on Hot Topics in Understanding Botnets, Berkeley, CA, USA: USENIX Association, 2007.
- [12] G. Gu, P. Porras, V. Yegneswaran, M. Fong, W. Lee, Bothunter: detecting malware infection through ids-driven dialog correlation, in: Proceedings of the 16th USENIX Security Symposium, August 2007. Online, available: <http://www.cyber-ta.org/releases/botHunter/>.
- [13] G. Gu, R. Perdisci, J. Zhang, W. Lee, Botminer: clustering analysis of network traffic for protocol- and structure-independent botnet detection, in: USENIX Security '08, 2008.
- [14] W.T. Strayer, R. Walsh, C. Livadas, D. Lapsley, Detecting botnets with tight command and control, in: Local Computer Networks, Proceedings 2006 31st IEEE Conference on, Nov.2006, pp. 195–202.
- [15] G. Starnberger, C. Krügel, E. Kirda, Overbot – a botnet protocol based on Kademlia, in: SecureComm 2008, 4th International Conference on Security and Privacy in: Communication Networks, September 22–25th 2008, Istanbul, Turkey, September 2008.
- [16] M. Allman, E. Blanton, V. Paxson, S. Shenker, Fighting coordinated attackers with cross-organizational information sharing, in: Hotnets 2006, 2006.
- [17] S. Katti, B. Krishnamurthy, D. Katabi, Collaborating against common enemies, in: IMC '05: Proceedings of the 5th ACM SIGCOMM conference on Internet measurement, ACM, New York, NY, USA, 2005, pp. 1–14.
- [18] DShield, Distributed intrusion detection system. Online, available: www.dshield.org, 2007.
- [19] Guofei Gu, Junjie Zhang, Wenke Lee, BotSniffer: Detecting Botnet Command and Control Channels in Network Traffic, in: The 15th Annual Network and Distributed System Security Symposium, 2008.
- [20] V.-H. Pham, M. Dacier, G. Urvoy Keller, T. En Najjary, The quest for multi-headed worms, in: DIMVA 2008, 5th Conference on Detection of Intrusions and Malware & Vulnerability Assessment, July 10–11th, 2008, Paris, France, July 2008.
- [21] N. Provos, A virtual honeypot framework, in: Proceedings of the 12th USENIX Security Symposium, August 2004, pp. 1–14.
- [22] C. Leita, V.H. Pham, O. Thonnard, E. Ramirez Silva, F. Pouget, E. Kirda, M. Dacier, The leurre.com project: collecting internet threats information using a worldwide distributed honeynet, in: 1st WOMBAT Workshop, April 21st–22nd, Amsterdam, The Netherlands, April 2008.
- [23] F. Pouget, M. Dacier, Honeypot-based forensics, in: AusCERT2004, AusCERT Asia Pacific Information technology Security Conference 2004, 23rd–27th May 2004, Brisbane, Australia, May 2004.
- [24] Z. Li, A. Goyal, Y. Chen, V. Paxson, Automating analysis of large-scale botnet probing events, in: ACM Symposium on Information, Computer and Communications Security, 2009.

- [25] Van-Hau Pham, Marc Dacier, Honey-pot traces forensics: the observation viewpoint matters, in: the 3rd Network and System Security, 2009.



Van-Hau Pham obtained his Bachelor degree in Computer Science from the University of Natural Sciences of Hochiminh City in 1998. He persuaded his Master degree in Computer Science from the Institut de la Francophonie pour l'Informatique (IFI) in Viet Nam from 2002 to 2004. Then he did his internship and worked as a full time research engineer in France for 2 years. He then persuaded his PhD thesis on network security under the direction of Professor Marc Dacier from 2005 to 2009. He is now lecturer at the International University of Hochiminh City. His main research interests include network security,

network protocols.



Marc Dacier joined Symantec as the director of Symantec Research Labs Europe in April 2008. From 2002 until 2008, he was a professor at EURECOM, France (www.eurecom.fr). He was also an associate professor at the University of Liege, in Belgium. From 1996 until 2002, he worked at IBM Research as the manager of the Global Security Analysis Lab. In 1998, he co-founded with K. Jackson the "Recent Advances on Intrusion Detection" Symposium (RAID). He is now chairing its steering committee. He is or has been involved in security-related European projects for more than 15 years (PDCS, PDCS-2, Cabernet,

MAFTIA, Resist, WOMBAT, FORWARD). He serves on the program committees of major security and dependability conferences and is a member of the steering committee of the "European Symposium on Research for Computer Security" (ESORICS). He was a member of the editorial board of the following journals: IEEE TDSC, ACM TISSEC and JIAS. His research interests include computer and network security, intrusion detection, network and system management. He is the author of numerous international publications and the holder of several patents.