Towards Optimal Sensor Placement Strategies for Early Warning Systems

Jan Göbel Philipp Trinius

Laboratory for Dependable Distributed Systems, University of Mannheim, Germany
{goebel|trinius}@informatik.uni-mannheim.de

Abstract: A network early warning system consists of several distributed sensors to detect malicious network activity. The effectiveness of such early warning systems critically depends on the sensor deployment strategy used. We therefore analysed attack patterns of malicious software collected at sensors worldwide to determine an optimal deployment strategy. Our results show that due to the small numbers of attackers shared among networks, the benefit of large-scale sensor deployment is rather limited. However, there is some evidence that world-wide geographical distribution of sensors has some beneficial effect on the average early warning time.

1 Introduction

Autonomous Spreading Malware and Early Warning: Autonomous spreading malware (malicious software) like worms or bots are one of the most dangerous threats in the Internet today. Once infected, hosts start to scan large network ranges for more vulnerable machines that can be exploited. Acting in a coordinated fashion, bots can launch denial of service attacks or initiate high-volume spam campaigns. The extreme dimensions of botnets (for example with approximately 350.000 machines [SGE+09]) make them a hard opponent for security systems. It is therefore necessary to develop ways to detect emerging threats as early as possible.

Many different network and host based security solutions have been developed in the past to counter the threat of autonomous spreading malware. Among the most common detection methods for such attacks are *honeypots*. Honeypots offer vulnerable services and collect all kinds of information about hosts that probe and exploit these services. Today, honeypots act as efficient sensors for the detection of attack activity in the Internet. Such sensor data can be combined with other network data to form a partial view of the status of a network.

Given an attack phenomenon that occurs in location L_1 at time T_1, it should be expected that the same phenomenon happens later at some time $T_2 > T_1$ in a different location L_2. If $T_2 - T_1$ is sufficiently large, the hope is that the network at location L_2 can "prepare" for the attack in time. For example, if some corporate network in Germany suffers from attacks by a particular new worm, the network providers could warn other organizations of an upcoming threat. Systems that implement this idea are often called *early warning*

systems. The *Internet Malware Analysis System* (InMAS) [EFG+09, EFG+10], developed at the University of Mannheim, is a prototype for a German early warning system for malware activities. It consists of several different sensor systems, including honeypots, to cover the most common attack vectors of current malware.

Effectiveness of Early Warning: The idea of early warning is rather intuitive. However, it is still unclear whether an early warning system for malware activities can in fact be effective. How much can an organization at location L_2 profit from the data collected at location L_1? This obviously depends very much on the number of sensors deployed and the IP address ranges they cover. It also critically depends on a sufficiently large period $T_2 - T_1$.

There exist some detailed investigations of the attack patterns of autonomous spreading malware [GHW07, BHB+09, SGL04] that investigate (1) what vulnerabilities are frequently exploited, (2) where attackers originate from, (3) what kind of malware is spreading, and (4) what the effects are to the infected hosts. In this paper, we abstract from particular exploits and malware and focus on *correlations* between attacks in different networks. This has direct effect on the effectiveness of early warning systems on a national or even global scale. Therefore, the results from our research can help to develop sensor deployment strategies for these scenarios.

Further research in this area focused on recorded connection data collected at the border gateway [APT07] of an academic network and on data gathered from blackhole deployments [CBM+04], i.e. unused address space. In contrast, our observations are based on data that (1) is clearly identified as being exploits of known vulnerabilities and (2) was collected from used address space, mainly of academic nature. Our work can therefore further substantiate previous research.

Contributions: To study the effectiveness of early warning systems, we first looked at the problem from a *national* perspective, i.e., we collected attack data from different honeypot installations across Germany. We investigated the question whether attack patterns at one location also occur at another and much time lies between such events. We then turned to a more *international* perspective and studied correlations between the German data and data from two external locations, one in China and the other one in Italy. We therefore refine the investigations and explore how sensors located further away from the German sensors influence the effectiveness of a national early warning system.

Overall, we collected attack data from a four month period in 2009. The sensors in Germany covered a total of 15.167 IP addresses, with the majority residing in a single /16 network range. This network range also served as the basis for studying attacker behavior on several adjacent /24 networks.

The major observations from our analysis are:

- We show that, within individual /24 networks, attackers mainly target low-numbered IP addresses, i.e., those between .1 and .127. This is independent both of the particular /24 network and the time of day. We also show that attackers launch attacks from a wide range of IP addresses and that the majority of such IP addresses is used for a small number of attacks only (Section 3).

- Within the low-numbered IP addresses, the attack probability decreases almost linearly from .1 to .127. Thus low-numbered IP addresses are a good choice for sensor deployment (Section 4).
- We show that the number of *shared* attackers over all networks is very small, i.e., the attacks launched from the same IP address almost never occur in many different /24 networks. There is no indication that geographical distance reduces the probability that two sensors will see an attack from the same IP address (Section 5).
- We show that in almost all observed attacks the reaction time, i.e., the time to warn others or roll out patches, is far too short, due to the low number of recurring attackers, and the fast scanning mechanisms used. There is some indication of an increase in average detection time if international sensors are used (Section 6).
- We observed and classified the scanning patterns of attackers. Most attackers performed random scanning. This observation can (partly) be used to explain the above findings (Section 8).

Overall, the results provide valuable information on sensor placement, but also show that early warning systems cannot (always) guarantee timely attack prediction.

2 Measurement Setup

Honeypot Sensors: For data collection we used two similar low-interaction honeypots, namely Nepenthes [BKD+06] and Amun [Goe09]. Although the main focus of both applications is to collect the binaries of autonomous spreading malware, they also generate log files that were the basis for our experiments.

Being low-interaction honeypots, our sensors were not designed to trick a human attacker; they just emulate exploitable services and allow no real interaction. In total the honeypots emulated 43 different services that represent a large percentage of the attack vectors of autonomous spreading malware today (see Appendix [Goe10] for a complete list). As there exist other ways for malware to propagate, our findings are restricted to such an automated setting only.

Data Points and Definitions: The data consists of individual data points that contain the following information: (1) geographical location of sensor, (2) IP address of sensor, (3) timestamp, (4) attacked port/service, and (5) IP address of attacker

Every individual data point represents a *successful exploit* of a service emulated by a honeypot. We use the terms *attack* and *exploit* synonymously in this paper. Note that exploits targeting multiple emulated services of the same honeypot count as multiple attacks. We identify an *attacker* with the source IP address of the attack. Similarly, we identify a sensor with its IP address. Networks are described as /24 and range from .1 to .255 (class C). We therefore use the term /24 and class C synonymously.

Geographical Distribution: For our measurement study we received data points from honeypots located at five different cities: Dresden (Germany), Aachen (Germany), Mannheim

(Germany), Milano (Italy), and Macau (China). Germany is covered best with three honeypot installations that consist of a total of 15.167 sensor IP addresses, the majority of these residing in a single /16 network range at Aachen. China and Italy both only have a single sensor IP address to record attack data.

Measurement Periods: The time period of our measurement study is split into two ranges. The first period lasted from April, 29, 2009 until May, 14, 2009 and the second from June, 10, 2009 until September, 14, 2009. In total, we have almost four months (113 days) of data. The only exception is the honeypot running in Macau (China), here the measurement period lasted from July, 1, 2009 until October, 27, 2009. In total we have an overlap of 76 days between the Chinese sensor and all other sensors.

3 Positioning of Sensors

Typically, only two IP addresses of a /24 network are not used for productive systems. The first address (.0), and the last address (.255), which is reserved for broadcast communication. The second address (.1) is usually assigned to the gateway for the network. Therefore, a common assumption is that an attacker scans the complete range between .1 and .254 for vulnerable hosts. To detect such an attacker, it would suffice to randomly deploy sensors across the network range.

To verify or even falsify this predication we aggregate the number of successful exploits recorded at each IP address of all monitored /24 networks of the Aachen honeypot installation, see Figure 1a. The figure shows the number of attackers seen at the individual IP address of each of the monitored /24 networks. There is a clear preference of low IP addresses and an interesting drop in the number of attackers at the center of each of the IP address ranges. Additionally, the number of attackers decreases slightly with the higher number of /24 networks.

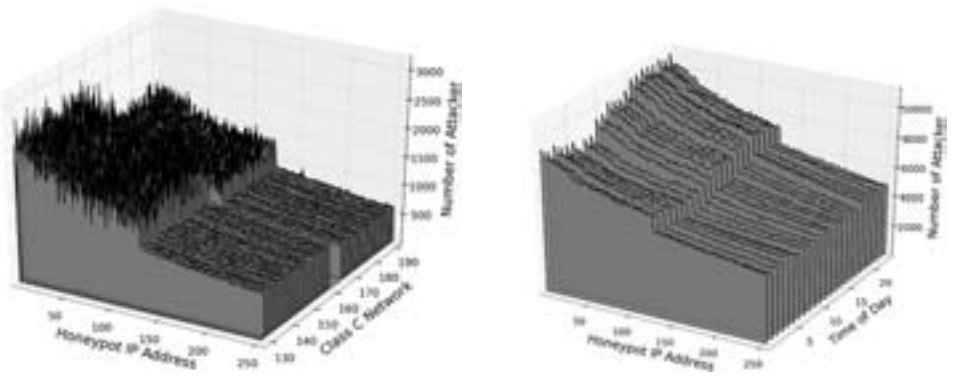

(a) Number of attackers per IP address in all /24 networks

(b) Number of attackers per IP address distributed over the day

Figure 1: Illustration of the number of attacker per IP address

As Figure 1a is an aggregate over the entire measurement period, we argued whether there is a dependence on the time of day. So we plotted the number of attackers again, this time for the different hours of a day (Figure 1b). Instead of generating a graph for each of the /24 networks, we merged them all into one and plotted them for each hour of a day. Figure 1b shows that there is an increase in the number of attackers for the late hours, however, the overall tendency we already showed remains.

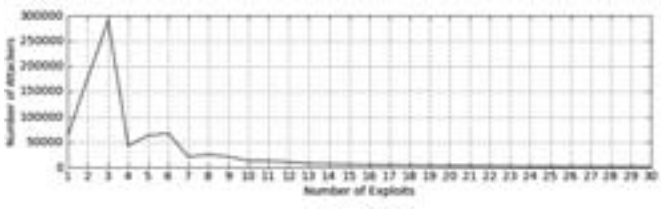

Figure 2: Number of exploits performed by attackers during the measurement period.

The reason for the uneven distribution of attacks within a network is that most attackers we monitored performed less than ten attacks during the complete measurement period. Figure 2 demonstrates the decrease of the number of exploits per attacker monitored during the complete measurement period for the Aachen honeypot installation. Note that for readability reasons the graph was cut at attackers performing more than thirty exploits. At the Aachen honeypots we monitored a total of 925.998 different attacking hosts and 83.64% have exploited the honeypot less than ten times. This phenomenon coincides with observations from previous work [GHW07].

Overall, using free IP addresses at the end of an address space as intrusion sensors makes less sense than placing sensors at specific points in between productive systems.

4 Number of Sensors

One way to measure the effectiveness of an early warning system is to count the number of attackers that are reported compared to the total number of attackers that targeted a network, i.e., the *detection ratio*. If we have the chance to place n sensors within a /24 network, clearly the most effective placement is to start with the IP address .1, then add .2, .3, up to $.n$. From the results of the previous section this strategy is even optimal.

Based on this observation we started with the lowest IP address as sensor and continuously added one address at a time (*linear increase*) until the complete network is covered with sensors. Figure 3a shows that a saturation of the detection ratio occurs in case 50% of the network is covered by sensors, starting at the lowest IP address of the network and increasing the number of sensors by one. Thus, if we deploy 127 sensors in the lower part of the /24 network, we achieve a detection ratio of 99.18%, i.e., our sensors detect 99.18% of all attackers that targeted this network during our measurement period.

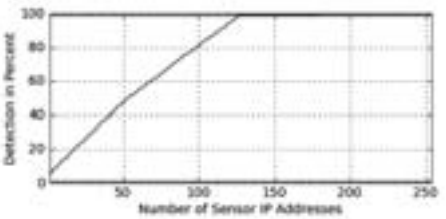

(a) Relation between the number of sensors and the detection ratio in a single /24 network.

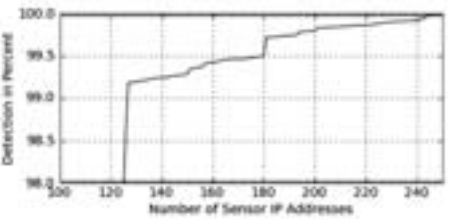

(b) Increase of detection ratio with more than half of the address space covered with sensors

Figure 3: Relation between the position of sensors and the detection ratio.

Figure 3b displays the upper part of Figure 3a starting at 98% detection ratio. The Figure illustrates the increase of the detection ratio when deploying more than 127 sensors.

5 (Globally) Shared Adversaries

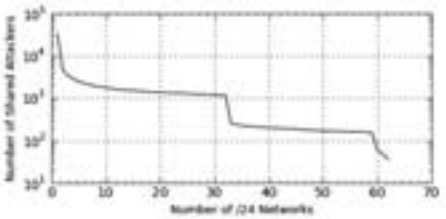

(a) Decrease in the number of shared attackers with an increase of monitored networks

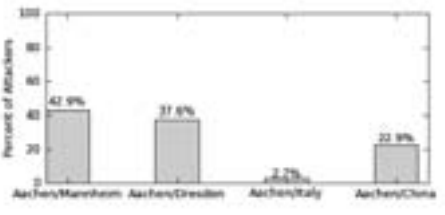

(b) Shared attackers among geographically distant locations

Figure 4: Shared attackers according to adjacent networks and distant location

We investigated the number of shared attackers with the increase of monitored /24 networks. Figure 4a illustrates the results and shows the clear decrease in shared attackers among all the /24 networks of the Aachen honeypot installation. In total we monitored 925.998 different attackers and end up with only 37 attackers that exploited systems in all networks.

Thus for an early warning system it does not suffice to have few sensors located in a single /16 network range. In order to collect information about as many attackers as possible it is required to have sensors in (almost) every /24 network.

Figure 4b further supports this point. The graph shows the percent of attackers shared between geographically distant sensors. We compared attackers seen at all our honeypot locations with the complete sensor network maintained in Aachen. The total number of attackers monitored at Mannheim, Dresden, Italy, and China are 577, 234, 4.202, and

24.471, respectively.

6 Attack Distribution and Detection Times

Another important factor with the detection of network incidents and the placement of sensors, is the time until the first detection of an attacker. The question is, how does the sensor deployment strategie affect the time of first detection for an attacker.

We first take a look at the 37 attackers that exploited honeypots in all /24 networks of the Aachen honeynet.

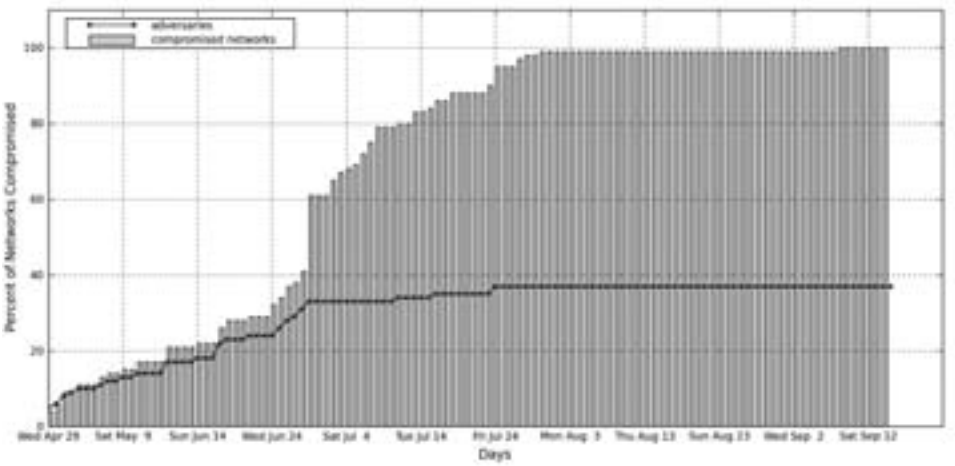

(a) Cumulative number of attacked /24 networks by shared attackers

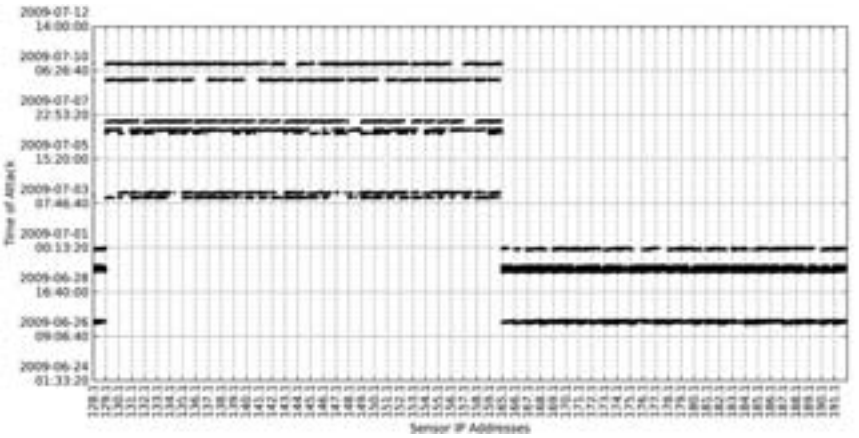

(b) Individual attacker exploiting hosts in all networks

Figure 5

Figure 5a shows the time needed for all 37 shared attackers to cover the complete network range. As not all attackers started to attack at the beginning of the measurement period, we indicated in the figure the cumulative number of attackers (up to the maximum of 37) that were present at a given time. For example, five attackers that scanned the whole network appeared on the first day of the measurement period. All 37 attackers appeared half way through the measurement period.

The figure shows that it took almost the complete measurement period until all networks had been attacked by all attackers. Thus for some attackers it is possible to react upon first detection and protect further networks. However, when looking at other attackers the detection time can be rather short.

We investigated the behavior of selected attackers further. Figure 5b illustrates the timeline of exploits performed by a single attacker out of the 37 mentioned above during our measurement period. Almost 50% of all /24 networks of the Aachen honeynet were exploited by this attacker on a single day. In this case, reaction time is rather short for the first half of the networks, but sufficiently large for the rest of the networks. This kind of parallel exploit behavior is common for 30 of the 37 attackers.

The fastest attacker exploited the complete Aachen honeynet within 1:33:58 hours. Although the attacker sequentially exploited host by host (see also Section 8), the time to react is extremely short even for all /24 networks. Appendix B provides a complete list of the first and last attack time and the used scanning mechanisms for all of the 37 shared attackers.

Investigating the attack times further, we could not observe any patterns indicating a particular order in which attacks occur. The 37 attackers that exploited honeypots in all /24 networks of the Aachen honeynet did not attack networks in the same order (starting for example with the lowest or highest number). Instead we noticed 22 different /24 networks at which the attacks were started. Furthermore, these networks seem to be randomly chosen.

7 Convenience of Geographical Distribution

After looking at attack times within and inbetween the /24 networks, we studied the correlation with attackers monitored at sensors that were geographically distributed.

Out of the 248 shared adversaries between Aachen and Mannheim, 64 exploited sensors in Mannheim first. The lowest time difference between two attacks of a shared attacker are 14 seconds, whereas the highest difference is 135 days.

The results with Dresden are similar. A total of 37 out of the 88 shared adversaries exploited honeypots in Dresden first, with a minimal time difference of 44 seconds and maximum of 128 days.

Out of the 114 monitored shared adversaries of the Italian sensor, 17 attacked Aachen after exploiting the Italian sensor. The fastest of these attackers still needed 6 minutes before reaching the Aachen sensors. The slowest needed 137 days.

Finally, for the Chinese sensor we detected 546 attackers out of the 5.605 shared adversaries that first exploited the honeypot installation in China and afterwards in Aachen. The shortest time to react upon such an attacker was 14 seconds whereas the longest time was 179 days.

L_1 / L_2	# Adv.	first $T_2 - T_1$	avg. $T_2 - T_1$
Mannheim/Aachen	64	120 hrs.	18 days
Dresden/Aachen	37	25 hrs.	9 days
Italy/Aachen	17	430 hrs.	45 days
China/Aachen	546	27 hrs.	44 days

Table 1: Summary of shared attackers, reaction time $T_2 - T_1$ for first shared attacker hitting remote location L_1, together with average reaction time over all shared attackers.

Table 1 summarizes our findings on shared attackers among geographically distributed sensors. The results show that having more distant sensors deployed helps to increase the average time to react upon an incident inflicted by shared attackers by at least 25 hrs. Furthermore, all sensors increase the total number of new, previously unseen adversaries that are then detected. Interestingly, the average delay correlates with geographical distance.

From the findings in this section we can conclude, that there is no clear geographical correlation between national sensor data, but there seems to be a correlation between national and international sensor data in the sense that international sensors have a substantially larger average reaction time for shared attackers.

8 Scanning mechanisms

During our measurement study we monitored over 955.476 unique attacker IP addresses. With many of these IP addresses showing up at several sensors, we were able the determine the scanning mechanisms used. Overall we can distinguish between four scanning mechanisms:

(1) random scanning (attacker selects /24 networks at random), (2) parallel scanning (attacker targets sensors from many /24 networks in parallel), (3) local sequential scanning (attacker targets sensors within /24 networks sequentially), and (4) global sequential scanning (attacker targets /24 networks sequentially). Table 2 shows the distribution among those four scanning mechanisms for 3.380 attackers which we observed at ten /24 networks of the Aachen honeynet.

Random scanning is the most primitive and difficult to detect scanning mechanisms. As Figure 6a shows, it is impossible to reason where and at which point in time, the attacker is seen again. Note that we picture ten different /24 networks, but the attacker appears only at a few sensors. Thus considering the notation from the introduction, it is not possible to determine L_2 and T_2 upon the observation of L_1 and T_1.

Scan Method	Percent
Random Scanning	55.6%
Sequential Scanning (local)	36.7%
Parallel Scanning	7.4 %
Sequential Scanning (global)	0.3 %

Table 2: Distribution of Scanning mechanisms across 3.380 shared attackers.

Parallel scanning refers to the scanning mechanism show in Figure 6b. The malware attacks sensors out of all /24 networks in parallel. For every /24 network it attacks several sensors within a very short time period. Even if an attacker who is using a parallel scanning mechanism is very easy to identify, the parallelism eliminates the time to react to almost zero, i.e. $T_2 - T_1$ is to low to react.

Local sequential scanning describes a very frequent scanning mechanism. The malware attacks several /24 network in parallel, but within each network it runs a sequential scan (see Figure 6c). Most of these scans are performed in increasing sequential order according to the IP addresses of the sensors. This scanning mechanism is good for exploiting a lot of systems without being noticed. The attacker can easily extend the gap between two attacks to stay below the threshold for intrusion detection systems. Figure 6d shows an attacker performing a slow scan.

In *global sequential scanning*, attackers perform a sequential scan on both: (1) all the /24 networks and (2) the IP addresses within each network. As shown in Figure 7, this scan is very obvious and can easily be identified. Due to the global sequential behavior it offers a relatively long time to react. In this case it is fairly simple to determine both L_2 and based on the speed of the scanning also T_2 of an attack. A classical example for malware that shows this kind of scanning behavior is the blaster worm [Naz03].

For the shared attackers of all /24 networks of the Aachen honeynet, we exclusively observed local sequential and parallel scanning mechanisms (see Appendix B for details). Approximately one-third is performing a sequential scanning, the remaining two-thirds scanned all the /24 networks in parallel. As Figure 5b already indicates, most of the attackers split the scanning process into two pieces and first scanned the higher range of /24 networks (159 to 191). In average the attackers exploited hosts of all networks for around 67 days. The fastest attacker was only observed for a single day, whereas the most persistent attacker was observed 137 days.

9 Conclusions

In this paper we investigated four month of honeypot attack data in order to determine sensor deployment strategies to achieve optimal effectiveness for early warning systems with focus on autonomous spreading malware.

In our view, the most important findings are:

(1) Using free IP addresses at the end of an address space as intrusion sensors makes less sense than placing sensors at specific points in between productive systems.

(2) It does not suffice to have few sensors located in a single /16 network range to achieve optimal attacker detection. It is rather required to have sensors in (almost) every /24 network.

(3) Adding more sensors to an early warning system can increase the time of reaction for those attackers that are shared. Especially international sensors have a substantially larger average reaction time for shared attackers. Therefore, even national early warning systems must deploy sensors at several distant locations, to achieve usable reaction times.

Acknowledgments

We would like to thank the Center for Computing and Communications of RWTH Aachen University, especially Jens Hektor, for their work on maintaining one of the biggest honeynet installations in Germany. We would like to thank Jens Syckor from TU Dresden, Matteo Cantoni from Italy, and Eric Chio from China for sharing their honeypot data. We also would like to thank Andreas Dewald for proof-reading and commenting on an earlier version of this paper.

References

[APT07] Mark Allman, Vern Paxson, and Jeff Terrell. A Brief History of Scanning. In *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement (IMC)*, New York, NY, USA, October 2007. ACM.

[BHB+09] Ulrich Bayer, Imam Habibi, Davide Balzarotti, Engin Kirda, and Christopher Kruegel. A view on current malware behavior. In *Proceedings of the 2nd Workshop on Large-Scale Exploits and Emergent Threats (LEET)*, Boston, USA, April 2009.

[BKD+06] Paul Baecher, Markus Koetter, Maximillian Dornseif, Thorsten Holz, and Felix Freiling. The nepenthes platform: An efficient approach to collect malware. In *Proceedings of the 9th International Symposium on Recent Advances in Intrusion Detection (RAID)*, pages 165–184. Springer, 2006.

[CBM+04] Evan Cooke, Michael Bailey, Z. Morley Mao, Danny Mcpherson, David Watson, and Farnam Jahanian. Toward understanding distributed blackhole placement. In *Proceedings of the 2004 ACM Workshop on Rapid Malcode (WORM)*, pages 54–64. ACM Press, 2004.

[EFG+09] Markus Engelberth, Felix C. Freiling, Jan Goebel, Christian Gorecki, Thorsten Holz, Philipp Trinius, and Carsten Willems. Frühe Warnung durch Beobachten und Verfolgen von bösartiger Software im Deutschen Internet: Das Internet-Malware-Analyse System (InMAS). In *11. Deutscher IT-Sicherheitskongress*, Bonn, Germany, May 2009.

[EFG+10] Markus Engelberth, Felix C. Freiling, Jan Göbel, Christian Gorecki, Thorsten Holz, Ralf Hund, Philipp Trinius, and Carsten Willems. The InMAS Approach. In *Proceedings of the 1st Workshop on Early Warning and Network Intelligence (EWNI)*, Hamburg, Germany, February 2010. https://eldorado.uni-dortmund.de/handle/2003/26689.

[GHW07] Jan Goebel, Thorsten Holz, and Carsten Willems. Measurement and Analysis of Autonomous Spreading Malware in a University Environment. In *Proceeding of 4th Conference on Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA)*, 2007.

[Goe09] Jan Goebel. Amun: A Python Honeypot. Technical report, Laboratory for Dependable Distributed Systems, University of Mannheim, 2009.

[Goe10] Jan Goebel. Amun: Automatic Capturing of Malicious Software. In *Proceedings of Sicherheit*, 2010.

[Naz03] Jose Nazario. The Blaster Worm: The view from 10,000 feet, August 2003. http://www.nanog.org/mtg-0310/pdf/nazario.pdf.

[SGE+09] Ben Stock, Jan Goebel, Markus Engelberth, Felix Freiling, and Thorsten Holz. Walowdac: Analysis of a Peer-to-Peer Botnet. In *Proceedings of the 5th European Conference on Computer Network Defense*, Milan, Italy, November 2009.

[SGL04] Stefan Saroiu, Steven D. Gribble, and Henry M. Levy. Measurement and Analysis of Spyware in a University Environment. In *Proceedings of the 1st Symposium on Networked Systems Design and Implementation (NSDI)*, pages 141–153, March 2004.

A Scanning Mechanisms

Figure 6 illustrates the four different types of scanning/exploiting mechanisms we detected during our investigation on the honeypot data. Figure 7 gives a more detailed view on the sequential scanning/exploiting.

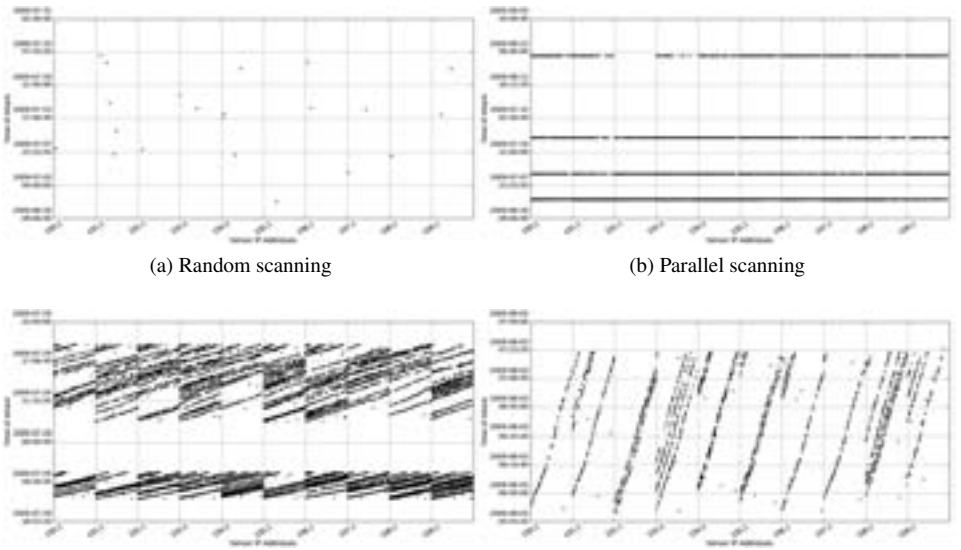

(a) Random scanning (b) Parallel scanning

(c) Sequential scanning (local) (d) Slow sequential scanning (local)

Figure 6: Scanning mechanisms used during attack sequences.

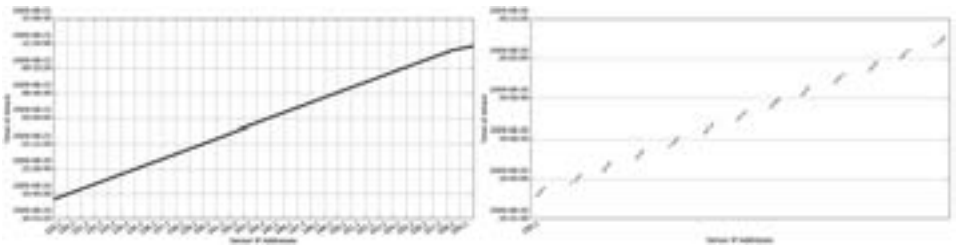

(a) Sequential scanning (global) (b) Details on a global sequential scan

Figure 7: Gobal sequential scanning.

B Common Attackers

Table 3 shows for all common attackers of all /24 networks of the Aachen honeynet the points in time the attacker was first and last seen. Additionally, the number of days during this two dates and the monitored scanning behavior is listed. The star (*) marks the fastest attacker, with approximately 1.5 hours to exploit hosts in all /24 networks of the Aachen honeynet. Moreover, it was the first scanning the complete honeynet during the monitored period.

Attacker	First seen	Last seen	Days	Scan
xxx.xxx.69.42	Apr 29 00:00:15	Sep 9 23:02:13	134	parallel
xxx.xxx.246.190	Apr 29 12:04:58	Aug 9 07:28:33	103	parallel
xxx.xxx.191.229	Apr 29 13:37:50	Sep 12 16:21:51	137	parallel
xxx.xxx.215.199	Apr 29 18:20:26	Sep 5 12:24:36	130	parallel
xxx.xxx.251.85	Apr 29 18:43:20	Aug 21 20:45:34	115	parallel
xxx.xxx.163.2	Apr 30 15:28:19	Jul 8 22:00:11	70	sequential
xxx.xxx.122.82	May 1 12:25:51	Sep 14 12:11:17	137	parallel
xxx.xxx.93.188	May 1 17:20:56	Sep 11 22:10:36	134	parallel
xxx.xxx.123.5	May 2 18:10:19	Sep 14 16:30:18	136	parallel
xxx.xxx.146.12	May 3 20:35:52	Jul 24 00:49:59	83	parallel
xxx.xxx.122.162	May 6 16:34:01	Aug 28 20:28:42	115	sequential
xxx.xxx.6.161	May 7 00:41:38	Jul 30 12:12:37	85	parallel
xxx.xxx.140.96	May 9 03:46:09	Sep 14 00:11:24	129	parallel
xxx.xxx.110.66	May 11 00:54:14	Aug 27 16:04:28	109	parallel
xxx.xxx.242.175	Jun 10 00:00:00	Sep 14 23:59:54	97	sequential
xxx.xxx.220.245	Jun 10 01:13:59	Aug 1 03:35:26	53	parallel
xxx.xxx.196.182	Jun 10 16:25:04	Aug 21 23:01:56	73	parallel
xxx.xxx.246.63	Jun 14 15:40:26	Sep 12 19:15:35	91	parallel
xxx.xxx.1.47	Jun 17 11:54:48	Sep 9 10:25:31	85	parallel
xxx.xxx.123.7	Jun 17 15:02:57	Sep 14 16:30:18	90	parallel
xxx.xxx.123.6	Jun 17 15:03:27	Sep 14 16:30:20	90	parallel
xxx.xxx.123.4	Jun 17 15:03:18	Sep 14 16:32:19	90	parallel
xxx.xxx.49.97	Jun 18 03:13:36	Jul 11 20:34:22	24	parallel
xxx.xxx.210.182	Jun 21 11:22:40	Jul 8 01:16:23	18	sequential
xxx.xxx.175.65	Jun 25 13:16:20	Sep 14 14:17:50	82	parallel
xxx.xxx.145.111	Jun 25 22:42:53	Jun 29 17:31:55	5	sequential
xxx.xxx.247.43	Jun 26 08:35:10	Jun 29 08:22:11	4	sequential
xxx.xxx.226.238	Jun 26 15:03:36	Jul 5 17:57:23	10	sequential
xxx.xxx.204.112	Jun 27 00:18:13	Jul 10 15:37:40	14	parallel
xxx.xxx.193.222	Jun 28 22:57:23	Jun 29 06:51:11	2	sequential
xxx.xxx.53.76 (*)	Jun 28 23:18:25	Jun 29 00:52:13	2	sequential
xxx.xxx.136.107	Jun 29 02:33:06	Jul 1 23:58:25	3	sequential
xxx.xxx.246.215	Jun 29 08:20:20	Jun 30 11:38:25	2	sequential
xxx.xxx.232.86	Jul 4 19:46:23	Jul 31 21:42:07	28	parallel
xxx.xxx.101.211	Jul 16 05:06:15	Jul 24 13:13:52	9	parallel
xxx.xxx.122.216	Jul 24 06:41:50	Jul 24 11:22:19	1	parallel
xxx.xxx.187.187	Jul 24 02:50:56	Jul 24 17:34:54	1	sequential

Table 3: Attack dates and scanning mechanisms of common attackers