Malicious fault characterization exploiting honeypot data

Corrado Leita, Olivier Thonnard
Institut Eurecom
France
leita@eurecom.fr,
olivier.thonnard@rma.ac.be

Eric Alata
LAAS
France
ealata@laas.fr

Marco Serafini
TU Darmstadt
Germany
marco@deeds.informatik.tu-darmstadt.de

Vladimir Stankovic
City University London
United Kingdom
ek274@soi.city.ac.uk

Jouni Viinikka
France Telecom R&D
France
jouni.viinikka@orange-ftgroup.com

Urko Zurutuza
Mondragon University
Spain
uzurutuza@eps.mondragon.edu

1 Introduction

ReSIST is a NoE that addresses the strategic objective "Towards a global dependability and security framework" of the European Union Work Programme, and responds to the stated "need for resilience, self-healing, dynamic content and volatile environments". In the context of the Resist Network of Excellence, funds were allocated to one-year sub-projects addressing some of the "dependability and security gaps" identified by the network activity. We introduce here the outcomes and deliverables expected by one of these sub-projects, named "Honeypots", started in January 2008 and with expected completion at the end of 2008.

In order to assess the resilience of a system it is important to take into consideration malicious faults. In the specific case of a computer system, one of the most important classes of malicious faults is linked to Internet attacks. The gap GA2 in [13] shows how the lack of unbiased, representative and useful data poses a serious limitation to any attempt of quantifying and characterizing the security threats to which a computer system may be exposed. Despite the considerable number of ongoing projects [17, 3, 4, 1, 6] aiming at collecting data about Internet malicious activities, the lack of "good" data with respect to this objective is an important gap of the current state of the art. Also, there is a lack of rigorous methodologies and models that can be used to extract relevant information and trends from the collected data.

In this project we want to address an interesting sub-problem of this gap: classify and characterize real-world Internet attacks linked to the spread of self-propagating malware. While the scenario of Internet attacks in the past was characterized by a few but extremely visible phenomena linked to the spread of worms, nowadays' scenario is far more complex. The objectives of the hackers have changed, and what started as a competition to gain visibility by infecting large populations of machines in very short time has now become a more organized covert activity aiming at economical profit. Internet is now dominated by a large number of different classes of malware aiming at silently taking control of computer systems in order to steal banking accounts, generate spam, steal credit card numbers, and many other profit-oriented operations. Even if the presence of these different classes is of public knowledge, we lack quantitative information on their spread, on their impact, on their evolution. This project aims at investigating the feasibility of using real data provided by honeypots to obtain such information.

2 Project activity

We plan to collect data taking advantage of a distributed honeypot deployment, developed at Institut Eurecom, called SGNET [8, 9]. SGNET aims at deploying small sensors in different locations of the IP space. These sensors emulate network services and, taking advantage of the ScriptGen technology [10, 11], achieve a very high level of verbosity at a very low cost. With respect to existing honeypot deployments, SGNET allows to achieve an extremely rich amount of information about the network activities hitting the honeypot, with special focus on code injection attacks. SGNET is in fact able to emulate the whole attack trace, understanding the presence of a successful code injection, and emulating its behavior ultimately downloading samples of the malware. This process allows to retrieve extremely valuable information about the intention of the attacker and its nature. This information is stored in a relational database for further analysis and constitutes the foundation over which we are planning to work within this project.

The "Honeypots" project aims at exploiting the information collected by the SGNET deployment. The project activity will follow two main lines of operation. Firstly, we will enrich the collected data through the comparison with other information sources. Secondly, we will dig into the collected data and try to abstract meaningful information taking advantage of existing algorithms developed by the participants in different contexts.

2.1 Correlation with other information sources

Some of the participants to the project manage different honeypot deployments taking advantage of different technologies. These technologies range from Honeynet Alliance deployments [1] to other high interaction techniques [2]. The reasons underneath the comparison of the information collected by different techniques are twofold. On the one hand, the diversity in collection techniques may allow to detect "blindness" of one approach to security events observed by the other approach. On the other hand, combining diverse information retrieved using different approaches potentially allows to enrich and consolidate the global view on the observed security events.

City University has deployed several high-interaction honeypots, which contain a mixture of Linux and Windows hosts and following the directions of the Honeynet Alliance [1]. These deployments emulate the traffic observed in corporate or SME (Small to Medium Business) environments. The network traffic collected in these honeynets will be compared with the SGNET data, e.g. exploratory analysis of the times between attacks on different hosts, operating systems, networks or geographical location will be detailed.

Mondragon University will proceed on the same line of City University, but from a different perspective. Mondragon University will provide a very rich dataset collected by around 10 honeypots of the Basque Honeynet Alliance [12]. This will allow to a very interesting comparison with the information provided by the SGNET deployment. The comparison will focus on different parameters than those analyzed by City University. For instance, both honeypot deployments being able to download malware using different strategies, we will compare the trends in malware download offered by the two technologies.

LAAS University will exploit the previous work on high interaction honeypots [2] to observe the behavior of attackers once they manage to get access to the core of the system. We plan to use this information to be able to provide answers to fundamental questions such as "what attackers do in the core of the system?" And "what are their objectives?". The goal is to better understand the motivations and methods of attackers. One way to achieve this goal is to study the sequence of commands performed by the attackers, using clustering techniques and symbolic time series.

2.2 Data enrichment through analysis

The information retrieved by the honeypot deployments is vast. In order to build characterizations of the observed security events we need a set of tools and techniques to mine the collected data and extract more meaningful aggregate information. We plan to take advantage of existing algorithms developed by some of the participants to the project to achieve this goal. The data mining techniques developed in the context of the Leurré.com honeypot deployment [5] can be of great help in organizing the collected data. Time series analysis developed in the context of IDS alert analysis [16] and statistical modelling [7] can help us in digging into the data filtering out uninteresting activities.

The data mining algorithms developed at Eurecom have been used so far to discover similarity patterns between the attacks over extended periods of time (e.g. several months or years). This long term analysis has already delivered interesting results about certain classes of malicious activities that seem to be correlated over long periods of time. In order to better address the highly dynamic and changing behavior new emerging threats, and also to understand how the threats evolve over time, we need to effectively identify the relevant periods of attack activities (i.e. the "attack events") on our sensors prior the execution of any correlative analysis. France Telecom R&D will provide its expertise in time series analysis [15, 14] to identify attack events within the honeypot time series. This can increase the quality and the significance of the results of the correlation algorithms developed by Eurecom. The final objective is to combine and automate both approaches (detection and correlation) so as to facilitate the processing of large quantities of data and to improve the discovery of meaningful information from honeypot data.

Eurecom will also provide an extension of its correlation framework in order to take advantage of the extensive information provided by SGNET. Such an extension will be used in the second half of the project to conduct an extensive analysis, so as to understand and characterize the evolution of malware activities in both the time domain and in the "spatial" domain (i.e. along the Internet IP space).

Attackers can use intruded nodes to disrupt the operation of distributed systems, for example by creating inconsistencies in a distributed file system where data is replicated on multiple machines to guarantee integrity and availability. Since tolerance to intrusions (or in general Byzantine faults) is expensive, it is important to estimate how many faults are to be tolerated. TU Darmstadt will look at the Honeypot data to answer this question for systems that are subject to untargeted attacks.

3 Conclusions

Summarizing, the "Honeypots" project aims at exploring the feasibility of characterizing Internet attacks by digging into the information collected by different honeypot deployments. We hope to achieve this through the joint efforts of researchers with different competencies and resources. The ultimate goal of the project is providing a significant step towards the collection of unbiased, representative and useful data on real world attacks.

Acknowledgment

This work is supported by the ReSIST Network of Excellence (contract number 026764).

References

[1] Know your enemy: GenII honeynets. *Know Your Enemy Whitepapers*, May 2005.

[2] E. Alata, V. Nicomette, M. Kaâniche, M. Dacier, and M. Herrb. Lessons learned from the deployment of a high-interaction honeypot. In *EDCC'06, 6th European Dependable Computing Conference*, Oct 2006.

[3] M. Bailey, E. Cooke, F. Jahanian, J. Nazario, and D. Watson. The internet motion sensor: A distributed blackhole monitoring system. In *12th Annual Network and Distributed System Security Symposium (NDSS)*, San Diego, February 2005.

[4] Caida Project. The UCSD Network Telescope, www.caida.org, 2007.

[5] M. Dacier, F. Pouget, and H. Debar. Leurre.com: On the advantages of deploying a large scale distributed honeypot platform. In *Proceedings of the E-Crime and Computer Conference 2005 (ECCE'05)*, Monaco, March 2005.

[6] DShield. Distributed Intrusion Detection System, www.dshield.org, 2007.

[7] M. Kaâniche, E. Alata, V. Nicomette, Y. Deswarte, and M. Dacier. Empirical analysis and statistical modeling of attack processes based on honeypots. *WEEDS 2006-workshop on empirical evaluation of dependability and security (in conjunction with the international conference on dependable systems and networks,(DSN2006)*, pages 119–124, 2006.

[8] C. Leita and M. Dacier. SGNET: a worldwide deployable framework to support the analysis of malware threat models. In *Proceedings of the 7th European Dependable Computing Conference (EDCC 2008)*, May 2008.

[9] C. Leita and M. Dacier. SGNET: Implementation Insights. In *IEEE/IFIP Network Operations and Management Symposium*, April 2008.

[10] C. Leita, M. Dacier, and F. Massicotte. Automatic handling of protocol dependencies and reaction to 0-day attacks with ScriptGen based honeypots. In *RAID 2006, 9th International Symposium on Recent Advances in Intrusion Detection, September 20-22, 2006, Hamburg, Germany - Also published as Lecture Notes in Computer Science Volume 4219/2006*, Sep 2006.

[11] C. Leita, K. Mermoud, and M. Dacier. Scriptgen: an automated script generation tool for honeyd. In *Proceedings of the 21st Annual Computer Security Applications Conference*, December 2005.

[12] Mondragon University. Euskalert, the basque honeynet alliance. Web page at `http://www.euskalert.net`, 2008.

[13] ReSIST NoE. From resilience-building to resilience-scaling technologies: Directions. *Deliverable D13*, 2007.

[14] J. Viinikka. *Intrusion Detection Alert Flow Processing Using Time Series Analysis Methods*. PhD thesis, University of Caen, Caen, France, Nov. 2006.

[15] J. Viinikka, H. Debar, L. Mé, A. Lehikoinen, and M. Tarvainen. Processing intrusion detection alert aggregates with time series modeling. *Information Fusion Journal*, 2008. Special Issue on Computer Security, to appear.

[16] J. Viinikka, H. Debar, L. Mé, and R. Séguier. Time series modeling for IDS alert management. *Proceedings of the 2006 ACM Symposium on Information, computer and communications security*, pages 102–113, 2006.

[17] D. Zamboni, J. Riordan, and Y. Duponchel. Building and deploying billy goat: a worm-detection system. In *18th Annual FIRST Conference*, June 2006.