Research Roadmap on Security Measurements

Xenofontas Dimitropoulos
ETH Zurich
fontas@tik.ee.ethz.ch

***Abstract*—In the context of the SysSec Network of Excellence call for consolidating the European (and international) systems security research community, this position paper aims at summarizing current research activities in the Communication Systems Group (CSG) of ETH Zurich relating to network security. Our research is aligned along three projects: 1) identifying, validating, and characterizing computer infections from intrusion alerts; 2) building privacy-preserving collaborative network security mechanisms based on efficient secure multiparty computation (MPC) primitives; and 3) dissecting Internet background radiation towards live networks with one-way flow classification. In addition, we highlight important directions for future research.**

I. Introduction

Our work is motivated by a number of inter-related problems. As modern malware rely increasingly on social engineering techniques, its becoming more important to develop reliable methods to detect with a small number of false positives systems within an organization that have been infected. Furthermore, to build better defenses its important to collaborate and share data about suspected attackers. Despite its clear advantages, collaboration is presently largely avoided due to privacy concerns.

Motivated by these problems, in the following sections we discuss our current and future research on security measurements. In Section II we first describe security datasets we have accumulated and use in our studies. Next, in Section III we describe our inference, validation, and characterization of computer infections. Sections IV and V outline our research on aggregating sensitive data using MPC and on classifying one-way flows, respectively. Finally, in Section VI we outline future directions and we conclude in Section VII.

II. Collecting and Correlating Diverse Types of Data about a Target Network

A fundamental problem with understanding security failures and with evaluating proposed defenses is the lack of security datasets from real-world environments. Datasets, like intrusion detection alerts and traffic traces, are vital for scientific research. In CSG, we have been collecting unsampled NetFlow data from the border routers of SWITCH since 2003 and have accumulated a rich archive of more than 70 Gbytes of compressed flow records. In addition, since 2009 we have been collecting Snort alerts from a sensor next to the edge router of the ETH campus. Using the Snort signature ruleset and the Emerging Threats (ET) ruleset, we observe on average three million alerts per day and have accumulated a rich archive of several million Snort alerts. *The collection of different types of data from a production network is essential for building a more complete picture of security incidents and for correlating footprints from different observation instruments*. Members of CSG are actively working on extending our archives with additional types of data. In particular, on-going efforts examine the possibility of collecting vulnerability reports from end-hosts, reports from a Deep Packet Inspection (DPI) system, and DNS data.

III. Identifying, Validating, and Characterizing Computer Infections from IDS Alerts

In collaboration with Elias Raftopoulos we are working on identifying, validating, and characterizing computer infections in a large academic network infrastructure from intrusion alerts. In particular, our work addresses the following problems (more information can be found in our technical report [1]):

Extrusion Detection from IDS Alerts: A fundamental problem of network intrusion detection systems is that they generate a large number of false positives. For example, *our IDS sensor generates on average three million alerts per day*. Given an archive of IDS alerts, an important problem is how an administrator can filter out the noise to identify actual security incidents. In this work, we are particularly interested in identifying computer infection within a monitored infrastructure, i.e., extrusion detection. Annotating a rich trace of IDS alerts with inferred security incidents is useful both for forensics investigations as well as for evaluating network defenses with realistic data. To address this problem, in our current research we have developed a novel alert correlation technique tailored for extrusion detection.

Validating Inferred Infections in a Production Network: Having inferred a number of suspected infections, it is very useful to thoroughly validate the incidents. This is particularly challenging in production environments, where validation might interfere with regular employee work. In this work, *we are interested in the challenging problem of remotely validating suspected infections on un-managed hosts within a production network*. In our current research, we are conducting a complex experiment in which an expert is assigned to manually validate live infections by collecting and analyzing data about the suspected host from a number of independent sources, including intrusion alerts, blacklists, vulnerability reports, host scanning, and search engine queries. The goal of manual assessment is to "connect the dots", i.e., correlate collected evidence and decide if they agree with the background knowledge about the suspected malware.

Characterizing Computer Infections: Infections are amongst the most critical events for computer administrators. Using our heuristic, we have identified several thousands infections that infected over a period of 9 months more than 11 thousand distinct hosts with static IP addresses. Given data about a large number of computer infections, it is important to systematically analyze infected hosts and derive insights for building better defenses. In our current research, we have characterized a number of different aspects of computer infections including infection times, types of infections, infected hosts, and observed spatio-temporal correlations. Among various interesting findings, we observe that the volume of inbound attacks to infected hosts increases rapidly after their infection and that infections exhibit significant spatial correlations, i.e., a new infection is more likely to occur close to already infected hosts.

IV. Building Privacy-Preserving Collaborative Network Security Mechanisms based on Secure Multi-party Computation

Internet security suffers from a fundamental imbalance: although attackers are globally spread and well coordinated, individual network domains are isolated to analyzing only local data, when in need to diagnose global security problems. If independent network domains collaborated, then it would be possible to design much more effective network monitoring and security mechanisms. For example, several collaborating networks can identify and blacklist spammers much faster and with higher accuracy than any single network alone. Despite its obvious advantages, collaboration is presently largely avoided because privacy laws, security policies, and competition prevent sharing sensitive data. To mitigate this problem we are working with Martin Burkhart on building efficient MPC primitives suited for collaborative network monitoring and security applications.

MPC appears as an ideal solution to the privacy-utility trade-off. On the one hand, any function can be turned into an MPC protocol and on the other hand the computation process provides strong privacy guarantees. Applying MPC on practical scenarios involving aggregation of network security data introduces the challenge of having to build very efficient protocols that can deal with voluminous input data. Our research spans a path starting from MPC theory, going to system design, performance evaluation, and ending with measurement. Along this path *a key contribution of our work is that we make an effort to bridge two very disparate worlds: MPC theory and network monitoring and security practices*.

In the theory front, we have designed optimized MPC comparison operations based on the observation that the performance of data-intensive MPC operations can be improved by not enforcing the widely-used constant-round paradigm. We have learned that constant-round protocols (at the cost of many multiplications) are not a panacea in MPC protocol design: *allowing many parallel invocations and removing the constant-round constraint enabled us to design protocols that substantially reduce the total run-time*.

In addition, we have designed four MPC protocols, namely the entropy, distinct count, event correlation, and top-k protocols. Our protocols have been inspired from specific network monitoring and security applications, but at the same time they are also general and can be used for other applications. For example, in our top-k protocol we have used sketches, an approximation data structure, to reduce expensive MPC comparison operations to much more efficient MPC additions and multiplications at the cost of a manageable approximation error [2].

We have implemented our basic operations and MPC protocols in the SEPIA library [3], [4], which we have made publicly-available. Our extensive performance evaluation shows that SEPIA operations are between 35 and several hundred times faster than those of existing comparable MPC frameworks. In addition, we have used our protocols with real-world traces from 17 customer networks of SWITCH to investigate the practical applicability of collaborative network monitoring and security based on MPC. We have investigated different ways the networks could have collaborated to troubleshoot an actual global network anomaly. This is the first work to apply MPC on real traffic traces and to demonstrate that *collaborative network monitoring and security based on MPC is both computationally feasible and useful for addressing real network problems*.

This work has led to one recent publication [3], while further research is carried out in the context of two projects. In the DEMONS FP7 collaborative European project, we are further optimizing the performance and extending the functionality of SEPIA, while in collaboration with IBM Research we have acquired support from the Swiss National Science Foundation to transfer the technology of the multiparty computation library to a privacy-preserving multidomain traffic flow analyzer.

V. Beyond Network Telescopes: One-way Flow Classification

Network telescopes extract Internet background radiation to unused IP address blocks and have been very useful in characterizing malicious patterns and events. An alternative way to study background radiation in edge networks is to find one-way flows, i.e., flows that never receive a reply. One-way flows are an important fraction of Internet traffic.

In our traces from SWITCH, *in 2004 two out of every three flows were one-way, while in 2008 one out of every three flows were one-way*. One-way flows provide a number of advantages for monitoring Internet background radiation. They can be easily extracted from live networks, can be used even when large unused IP address blocks are not available, and require fewer resources for instrumentation. To enable administrators and researchers to extract and analyze background radiation from one-way flows, we have been working with Eduard Glatz on *novel techniques for classifying one-way flows in a set of malicious and benign classes*. Our classification associates each one-way flows with up to 17 different signs derived from flow-level data and uses a set of expert rules to determine the appropriate class of an one-way flow. It does not require training (plug & play), is easily extensible, and is based on comprehensible classification rules. *We have used our classification to analyze a massive dataset of flow records summarizing 7.73 petabytes of traffic between 2004 and 2010* that crossed the border routers of SWITCH. In addition, we work on characterizing how the volume of background radiation changed over time and what is the persistence of local port numbers as top targets. Among our findings, we observe that the relative volume of background radiation towards the target network decreased sharply between 2004 and 2007 by 73% and remains almost constant since then. In addition, filtering the top-5 or 10 port numbers, reduces the volume of background radiation by 35.6% and 45.0%, respectively. We believe that one-way flow classification opens a number of new directions on exploiting background radiation for building better intrusion detection systems. More information can be found in our technical report [5].

VI. Future Directions

Based on our experience with our current research, we believe that the following open problems deserve further attention in the future:

1) **Assessing the privacy risk of the MPC output:** Although MPC based on Shamir's secret sharing scheme guarantees that the computation provides information theoretic privacy, the output of the computation may still pose privacy risks especially if an adversary can correlate it with background knowledge. Intuitively, aggregating data from multiple parties obscures the sensitive input of any individual party. However, it is still an open question *how to assess the privacy risk of the output of an evaluated MPC function and if the output needs to be perturbed before publication* to provide some rigorous privacy guarantees.
2) **Automating security assessment:** Our experience with validating a number of inferred computer infections taught us that *security assessment in large organizations is a magic art rather than a scientific process*. It requires collecting relevant evidence, leveraging the background knowledge of domain experts about the infrastructure and the suspected malware, and relying on an ad-hoc cognitive process to diagnose a suspected security incident. In the future, the steps of a security assessment process need to be supported by automated techniques that expedite or even replace the manual work of a security administrator.
3) **Correlating flow data with intrusion alerts:** Past research has extensively studied how an administrator can use intrusion alerts or flow data to detect attacks and other security incidents. Although these data sources have been studied in isolation, an important question is *how to best combine flow data with IDS alerts to detect security incidents more accurately*. Flow data provide more information about the normal behavior of a host, while intrusion alerts provide more details about a suspected incident. Their proper combination can lead to reducing the false positive rate of IDSs and to enriching flow-based anomaly detectors with more information about a suspected anomaly.

VII. Conclusions

This short paper aims at summarizing the current research of part of the Communication Systems Group of ETH Zurich that relates to security measurements. In addition, we pinpoint important future research directions. Our focus areas are on: 1) identifying, validating, and characterizing computer infections from intrusion alerts; 2) MPC-based collaborative network monitoring and security applications; and 3) one-way flow classification. More details about the summarized works can be found in [1], [4], [3], [5].

References

[1] E. Raftopoulos and X. Dimitropoulos, "Detecting, validating and characterizing computer infections from IDS alerts," ETH Zurich, TIK-Report 337, June 2011.

[2] M. Burkhart and X. Dimitropoulos, "Fast privacy-preserving top-k queries using secret sharing," in *International Conference on Computer Communications and Networks (ICCCN)*, 2010.

[3] M. Burkhart, M. Strasser, D. Many, and X. Dimitropoulos, "SEPIA: Privacy-Preserving Aggregation of Multi-Domain Network Events and Statistics," in *USENIX Security Symposium*, 2010.

[4] M. Burkhart and X. Dimitropoulos, "Privacy-preserving distributed network troubleshooting – bridging the gap between theory and practice," *ACM Transactions on Information and Systems Security (under submission)*, 2011.

[5] E. Glatz and X. Dimitropoulos, "Beyond network telescopes: One-way traffic classification," ETH Zurich, TIK-Report 336, June 2011.