METIS: a Two-Tier Intrusion Detection System for Advanced Metering Infrastructures

Vincenzo Gulisano Magnus Almgren Marina Papatriantafilou

Chalmers University of Technology
{vinmas,almgren,ptrianta}@chalmers.se

ABSTRACT

Specification-based intrusion detection systems, the main defense mechanism proposed so far for Advanced Metering Infrastructures, do not provide a comprehensive protection against the wide spectrum of possible attack scenarios. Challenging aspects in this context include the need for timely detection and for novel attack scenario modeling techniques.

This paper introduces *METIS*, a novel two-tier anomaly-based intrusion detection framework that targets such challenges. The framework provides a continuous and fully distributed processing of network traffic by relying on the data streaming processing paradigm. Attack scenarios can be specified by means of the traffic features they affect and their resulting patterns of malicious activities. We overview the framework, presenting the novel detection technique, and provide results from a case study.

Categories and Subject Descriptors

C.3 [**Computer Systems Organization**]: Special-Purpose and Application-Based Systems; C.2.4 [**Computer Systems Organization**]: Computer-Communication Networks—*Distributed Systems*

Keywords

Advanced Metering Infrastructure; Intrusion Detection; Data Streaming

1. INTRODUCTION

Advanced Metering Infrastructures (AMIs) are composed of communication-enabled metering devices that share data and are remotely controlled by energy utilities. Despite the limited number of real attacks documented so far, a considerable amount of possible attack vectors has been uncovered (e.g., by means of penetration testing). Even though specification-based Intrusion Detection Systems[1] (IDSs), the main defense mechanism proposed so far, may detect some attacks, they are not suited to provide a comprehensive protection against all possible attack scenarios, especially considering the manual labor required by a security expert to tune such systems to specific installations. More concretely, they might fail in detecting malicious activity residing in between the boundaries of a device's normal behavior (e.g., they might detect a malicious firmware that suppresses consumption readings but probably miss a firmware that lowers bills by slightly reducing reported consumption readings).

Challenges and contributions. We identify two main challenges for an IDS used in the context of AMIs. As discussed in [1], the AMI network consists of several independent networks and a centralized IDS would not have access to the overall traffic. Moreover, its processing capacity could easily be exhausted by the big, fluctuating volume of data generated by AMIs' devices. For these reasons, the IDS should (1) process data in a fully distributed and parallel manner in order to detect malicious activity timely and (2) avoid expensive per-site customization, by providing an efficient way to specify how malicious activities should be detected.

We present *METIS*[1], an anomaly-based IDS that provides continuous and fully distributed analysis of AMIs traffic. It relies on a novel two-tier detection method in which general attack scenarios are specified by means of (i) the traffic features they involve (and their inter-dependencies) and (ii) how the resulting patterns of suspicious events should be interpreted. We overview the framework, presenting its processing model and its detection technique, showing its use and also providing results from a case study.

2. METIS - OVERVIEW

Continuous and distributed traffic analysis. We aim at the design of an IDS that processes data in an online fashion in order to detect malicious activities timely. The need for continuous processing is also motivated by the evolving nature of AMIs, where new devices are deployed on a daily basis. We identify *data streaming*[4] as a suitable processing paradigm for the analysis of AMIs' data flows. With data streaming, on-line analysis is performed by means of directed acyclic graphs of operators referred to as *Continuous Queries*. Operators provide functionality such as data filtering, aggregation and joining. Data streaming allows for

[1] Named after the mythology figure standing for good counsel, advice, planning, cunning, craftiness, and wisdom.

e-Energy'14, June 11–13, 2014, Cambridge, UK.
ACM 978-1-4503-2819-7/14/06.
http://dx.doi.org/10.1145/2602044.2602072.

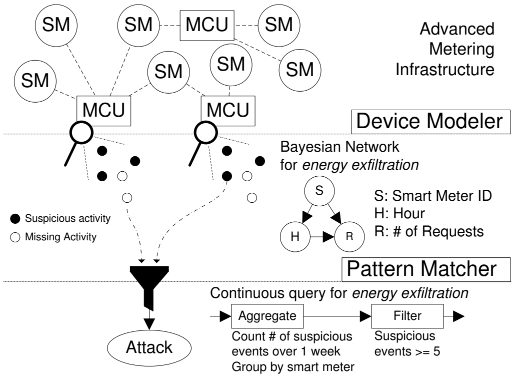

Figure 1: AMI architecture and *METIS* processing overview.

distributed and parallel processing that can scale accordingly to (i) the increasing number of devices deployed in the network and (ii) the increasing number of attack scenarios being monitored.

Two-tier detection. One of the critical aspects of a defense framework is the way the possible adversary goals need to be specified. To this end, we introduce a novel two-tier detection method in which attack scenarios are specified by the device's traffic features they affect and by the pattern of suspicious events they produce.

Figure 1 presents how the traffic analysis is carried out for a sample AMI composed by Smart Meters (SMs) and Meter Concentrator Units (MCUs) in charge of collecting energy consumption readings. Briefly speaking, *METIS* is composed of two main modules. The *Device Modeler* analyzes each device's traffic and spots suspicious or missing events applying an outlier detection method. The *Device Modeler* relies on a lightweight, yet efficient analysis technique, in which the traffic features that characterize the attack are expressed by means of a Bayesian Network. The secondary analysis is performed by the *Pattern Matcher*, which processes such alerts and requires the user to specify which patterns should be classified as an attack. Such analysis could be performed at dedicated servers, thus relying on more expensive analysis and correlation of the observed alerts.

3. CASE STUDY

Fine grained consumption readings reveal detailed information about household activities [3]. In an energy exfiltration attack scenario, an attacker that gains admin access to an MCU, requests a given SM to report its consumption readings. Notice that a specification-based IDS might fail in detecting such malicious activity if the MCU is actually allowed to retrieve such data from the target SM.

We monitor the traffic of approximately $1,000$ SMs connected to 40 MCUs during the months of September to December 2012. We inject 40 simulated energy exfiltration attacks (introducing unauthorized consumption readings) during the month of December. The *Device Modeler*'s Bayesian network and the *Pattern Matcher*'s continuous query are shown in Figure 1. The traffic features we take into account are the smart meter ID (S), the hour (H) and the number of requests (R) observed for each MCU. The *Pattern matcher* looks for MCUs reporting at least five suspicious requests during a period of one week.

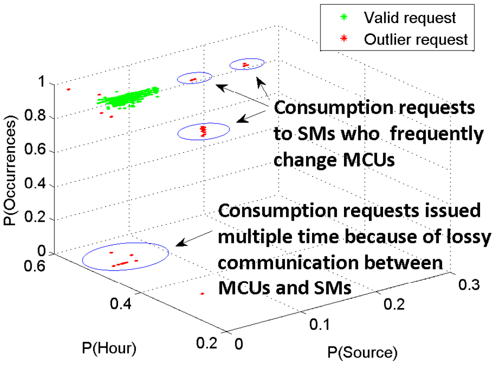

Figure 2: Local Outlier Factor applied to the Bayesian network probabilities.

From our evaluation, the framework is able to detect most of the injected attacks. Figure 2 presents the result of the outlier detection method (Local Outlier Factor [2]) and shows which consumption readings are marked as suspicious. Approximately 8% of the events observed at the MCUs are classified as suspicious by the *Device Modeler* ($4,000$ out of $50,000$), while the *Pattern Matcher* provides a True Positive Rate of approximately 91% (36 attacks are detected).

4. CONCLUSION

We propose a two-tier, anomaly-based intrusion detection framework that deals with the security and induced computational challenges of AMIs. We also present a case study that shows it can be flexible for the user to use and can achieve high detection rates. We plan to extend our research considering the detection of a variety of scenarios, including those whose detection is only possible through distributed evidence.

5. REFERENCES

[1] R. Berthier and W. H. Sanders. Specification-based intrusion detection for advanced metering infrastructures. In *IEEE 17th Pacific Rim International Symposium on Dependable Computing (PRDC), 2011*, 2011.

[2] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander. Lof: identifying density-based local outliers. In *ACM Sigmod Record*, 2000.

[3] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin. Private memoirs of a smart meter. In *Proceedings of the 2nd ACM workshop on embedded sensing systems for energy-efficiency in building*, 2010.

[4] M. Stonebraker, U. Çetintemel, and S. Zdonik. The 8 requirements of real-time stream processing. *SIGMOD Rec.*, 2005.

Acknowledgments

This work has been partially supported by the European Comission Seventh Framework Programme (FP7/2007-2013) through the SysSec Project, under grant agreement 257007, through the FP7-SEC-285477-CRISALIS project and through the collaboration framework of Chalmers Energy Area of Advance.