Collaborative Intrusion Detection in Smart Energy Grids

Emmanouil Vasilomanolakis, Mathias Fischer, Max Mühlhäuser
Telecooperation Group, Technische Universität Darmstadt,
Center for Advanced Security Research Darmstadt (CASED)
{manolis,mathias.fischer}@cased.de, max@informatik.tu-darmstadt.de

Peter Ebinger, Panayotis Kikiras, Sebastian Schmerl
AGT Group (R&D) GmbH
{pebinger, pkikiras, sschmerl}@agtinternational.com

The ongoing convergence of Industrial Control Systems (ICSs) with the Internet introduces many challenges from security perspective. Particularly, the smart energy grid as large ICS and critical infrastructure, requires especial protection as the consequences of its failure can be severe. However, even a careful system design cannot prevent all attacks in advance. For this reason, the smart grid requires an additional line of defence that can be provided by a Collaborative Intrusion Detection System (CIDS) to detect unknown and ongoing attacks. In this paper, we describe the requirements to a CIDS for deployment in the smart grid. Furthermore, we discuss the design choices for such a system and summarize the arising challenges in the deployment of CIDSs in smart grids as well as present initial ideas to address them.

Smart Grid, Collaborative Intrusion Detection, Monitoring

1. INTRODUCTION

In the past, Industrial Control Systems were isolated by design and thus built with no security concerns. Nowadays, the increasing interconnectivity needs of corporations enforce the convergence of ICSs with business networks that are connected to the Internet, which results in an increased vulnerability of ICSs. The next generation electrical grid, the *smart grid*, is a large ICS and a critical infrastructure. Fig. 1 shows a high-level view of a smart grid, which is a complex network consisting of energy providers, transmission networks, distribution networks, consumers but also prosumers that consume and produce energy at the same time. The implications in case of failures of energy networks can be severe, so that they need to be especially safe-guarded. However, even a careful system design and development cannot prevent all attacks. This is shown by recent attacks on ICSs, e.g., via Stuxnet, and requires advanced protection mechanisms. Hence, as an additional line of defence and to detect unknown attacks, intrusion detection and monitoring within the smart grid is required. This demands for the deployment of scalable monitoring mechanisms in all parts of the smart grid, as indicated by the monitors *(M)* in Fig. 1.

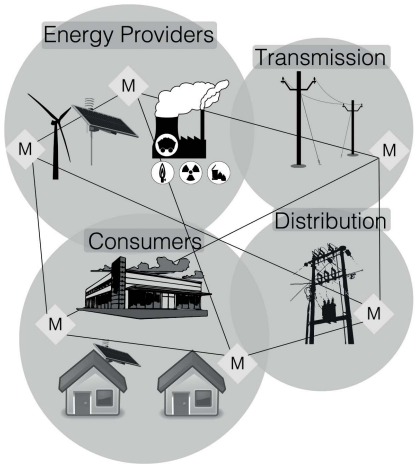

Figure 1: *Intrusion Detection in the Smart Grid via a set of Monitors (M).*

Intrusion Detection Systems (IDSs) have been extensively studied over the past years, e.g., Barry and Chan (2010). In comparison to isolated IDSs schemes, in which an isolated monitoring point observes a certain part of a network, more collaborative concepts have emerged. CIDSs and the arising challenges of their deployment are discussed in Zhou et al. (2010). Since the smart grid

is a large network and contains several components, as seen in Fig. 1, a scalable CIDS has to be applied to pair the different monitoring points.

CIDSs can be classified according to their architecture in centralized, decentralized, and distributed CIDSs. *Centralized* CIDSs consist of a number of monitoring points that pass their monitored data to a central analysis unit. In contrast, *decentralized* CIDSs (Fig. 2 on the left) utilize a hierarchical structure of monitoring points in which a preprocessing and correlation of the monitored data from layer to layer takes place. This continues until the aggregated data is finally processed by a central analysis unit. In *distributed* CIDSs (Fig. 2 on the right) the tasks of the central analysis unit are completely distributed onto the monitoring points.

The different CIDS architectures and their advantages and disadvantages with respect to smart grids are discussed in Section 3. Only a few recent proposals discuss directly intrusion detection in smart grid environments, e.g., Berthier et al. (2010) and Zhang et al. (2011). However, they focus on the detection level rather than on the architecture level and the challenges of intrusion detection in smart grids.

The contribution of this paper is twofold: In Section 2 we define the necessary requirements for the deployment of a CIDS in a smart grid. In Section 3 we present the available architectures and address the arising challenges in building a CIDS for smart grids. These challenges are related with the architecture of the CIDS, the distribution of monitored data, and possible data correlation as well as aggregation techniques that have to be applied. The main contribution of this paper is a summary of arising research problems when bringing CIDSs to smart grid environments and initial ideas to address them. Finally, Section 4 summarizes the paper.

2. REQUIREMENTS FOR SMART GRID IDSs

In this section, we describe basic requirements for IDSs intended for the protection of the smart grid.

Accuracy is mainly determined by the percentage of successfully detected attacks and the corresponding number of not detected attacks (false negatives). To rate the accuracy of an IDS also the number of falsely triggered alarms (false positives), has to be taken into account. Furthermore, for smart grids besides the monitoring of communication flows also a monitoring of energy flows is required. Moreover, a smart grid IDS should offer real-time detection of attacks to enable fast response and restoration strategies.

Overhead arises in terms of *computation* and *communication*. The techniques used to produce, collect, or correlate intrusion alerts, must have low computational overhead. Furthermore, the signaling inside the CIDS needs to be minimal due to possible network limitations, e.g., due to GSM/GPRS links.

Scalability requires that the performance of the IDS increases linearly with the size of the resources added, so that networks of arbitrary size can be protected. For that, the IDS should contain no bottlenecks and Single Point of Failures (SPoFs).

Resilience refers to the ability of the CIDS to maintain its availability, by providing an acceptable accuracy, in the presence of failures and attacks on monitoring points. Thus, this requires to prevent SPoFs and to provide graceful degradation abilities as well as fast restoration mechanisms.

Privacy and Confidentiality are important for a collaborative environment as the alerts that are exchanged may include private information that needs to be especially protected and should not be disclosed to all IDS components. Particularly, in a smart grid environment monitored data that crosses domain borders, poses the need for selectively forwarding information to protect privacy and confidentiality of business information.

3. DESIGN CONSIDERATIONS FOR CIDSs IN SMART GRIDS

To effectively design a CIDS for the smart grid, certain considerations have to be made. First, the optimal architecture for the CIDS needs to determined, which can be either centralized. hierarchical, distributed or a hybrid architecture. Second, the overlay ID space can be either structured or unstructured depending on whether efficient search functions or grouping of arbitrary peers is required. Finally, the data exchange within the CIDS has to be considered. especially which parts of the monitored data are to be shared internally.

CIDS Architectures - A CIDS for smart grids can either follow a centralized, decentralized, or distributed structure. Each of the available architectures exhibits certain advantages and disadvantages. A centralized CIDS achieves the highest accuracy rate as all the monitoring data is sent to a central analysis unit. However, these systems are not scalable. Hence, to build a scalable CIDS for the smart grid, an hierarchical or distributed approach (Fig. 2) has to be deployed. As there is no central analysis unit that is in possession of all the data, these systems do not offer accuracy rates equivalent to centralized

IDSs. Hence, the main research goal here is to deal with this trade-off between scalability and accuracy.

At first sight, an hierarchical IDS approach seems to be most suitable for protecting smart grids that are also hierarchically organized. However, hierarchical approaches suffer from certain disadvantages. Since data is aggregated bottom-up from layer to layer and in each layer information is lost, this class is maybe less accurate when facing sophisticated attacks, e.g., covert channel attacks (Zander and Armitage (2007)). Failures of single monitors in the hierarchical monitoring overlay can be compensated by self-configuration and repair mechanisms. However, the root node in the architecture remains a SPoF.

Alternatively, a distributed and flat IDS architecture can be applied to overcome the bottleneck and SPoF in a hierarchical approach, i.e., the root node. However, in such a flat architecture and without a central point, no monitor has the complete overview about the state of the whole CIDS. In contrast to a hierarchical system, in which information flows are directed from bottom-up towards the root node, they can be arbitrarily directed in distributed CIDSs. Hence, certain data may only be shared in between a subset of all monitors. Therefore, a special membership management within distributed CIDS is required that connects appropriate monitoring peers with each other, e.g., monitors that experience similar traffic patterns.

Neither hierarchical nor distributed CIDSs seem to fulfill all requirements for a deployment in smart grids. In hierarchical CIDSs the data aggregation amongst different layers of the hierarchy results in the loss of data per traversed layer, but provides a global view on the whole system at the root of the hierarchy. Distributed CIDS are much more flexible than hierarchical approaches, but provide no global view on the system anymore. Instead certain data may only be shared among a subset of all monitors. This has the advantage that data necessary for the detection of sophisticated attacks, e.g., covert channels, may not be lost while being aggregated via different layers as in hierarchical approaches.

Thus, we argue that a hybrid architecture that combines the flexibility of distributed approaches with the ability of having a central view of the monitored system is most beneficial. For that, the existing hierarchical structure of the smart grid could be utilized to build up a hierarchical organized monitoring overlay. At the same time, additional links amongst monitoring points in different layers can be established, to provide more flexibility. This results in the following question:

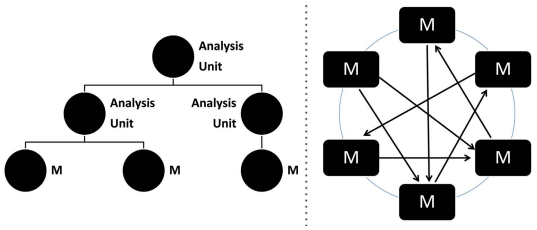

Figure 2: *A (left) decentralized, and (right) distributed IDS architecture.*

How can the advantages of hierarchical and distributed architectures be combined efficiently?

Structured vs. Unstructured Overlay - The ID space of a CIDS overlay can be either structured or unstructured. The basic advantage of a structured overlay, e.g., a Distributed Hash Table (DHT), is its guaranteed broadcast and search functionality. However, DHTs only allow single-attribute lookups, while complex multi-attribute searches are not feasible. Thus, it is essential to select the right property as key for the DHT. Moreover, in CIDSs besides distributing alert data immediately, an efficient storing of data is required. For this, DHTs cannot provide strict locality properties and the monitored data may be stored anywhere in the overlay network. This could violate the privacy requirement. However, a CIDS monitoring several transmission and distribution networks from different providers requires privacy. In contrast to structured architectures that enforce a strict ordering of peers according to their ID, unstructured approaches are more flexible in establishing overlay links. In a CIDS this can be exploited to establish overlay links based upon specific properties, e.g., commonly monitored traffic patterns. However, lookup functionalities and data sharing can only be realized via flooding that creates significant signaling overhead.

For this reason, we propose to establish communities (or clusters) within unstructured overlays, in which nodes exchange detailed data. These communities are established on the basis of specific properties, like similar traffic patterns. Monitoring points that observe similar traffic join the same community and exchange more detailed data among each other than with the rest of the overlay. As a result, even sophisticated attacks, e.g., distributed scans slowly manifested over time, can be detected. Moreover, locality can also be used as property to form communities of topological close monitoring points, e.g., to prevent information flows from one energy provider to the other. To establish these communities, coarse-grained and aggregated monitored data can be distributed in the whole network to spread information about traffic patterns via a probabilistic flooding protocol. All the peers that receive coarse-grained data

can form a community to exchange more detailed data. In this way, communities can operate jointly to create a more holistic view of a security incident and thus provide better accuracy in detecting attacks. Hence, the next question to be answered is:

How accurate is this community-based approach in detecting attacks, compared to existing ones?

Data Correlation - Depending on the attacker model and the overlay used, the information that is distributed can be coarse-grained, e.g., containing only IP-addresses of attackers to perform blacklisting, or the information can be more detailed, e.g., IP addresses and port numbers or complete packets. Other challenges that arise regarding the sharing of monitored data are related to storage issues. For instance, in a P2P based IDS it has to be determined which parts of the alert data should be stored locally and which should be shared amongst the peers.

Several correlation techniques have been proposed. For instance, Debar and Wespi (2001) proposed algorithms for correlating similar alerts based on simple properties, like bringing together identical alerts from different monitors. While these proposals are good in terms of overhead they reduce the system's accuracy. Other approaches, like Dain and Cunningham (2001), correlate alerts based on certain attack scenarios to detect more complex, but known, attacks. Furthermore, based on the multiple steps an attacker has to perform to successfully penetrate a system several techniques have been suggested, e.g., Cuppens and Miège (2002) and Cheung et al. (2003).

Since we described a hybrid architecture with both distributed and hierarchical components, the correlation can differ depending on the architecture. For instance, in the proposed distributed approach, correlation can be more sophisticated, while in the hierarchical approach it can be more simplistic, e.g., fusion of identical alerts from different monitors. These research challenges can be summarized with the following questions:

Which monitored data should be distributed and/or stored? What kind of correlation and aggregation mechanisms should be deployed and at which extent?

4. CONCLUSION

Since smart grids are a critical infrastructure, their resilience against attacks and failures is of high importance. As an additional line of defense, scalable monitoring mechanisms for intrusion detection are required. Hence, a CIDS with cooperating monitoring points is needed. In this paper, we discuss the challenges of deploying CIDSs in smart grids. Moreover, we offer some initial ideas on how to overcome problems related to the architecture level and to determine the sort of information that should be exchanged amongst the monitoring points. We propose a hybrid approach by combining a hierarchical architecture with a distributed one. This allows to establish *communities* for the exchange of more detailed information.

REFERENCES

Bazara I A Barry and H Anthony Chan. Intrusion Detection Systems. In *Handbook of Information and Communication Security*, pages 193–205. Springer Berlin, 2010.

Robin Berthier, William H. Sanders, and Himanshu Khurana. Intrusion Detection for Advanced Metering Infrastructures: Requirements and Architectural Directions. In *IEEE SmartGridComm*, pages 350–355, 2010.

S Cheung, U Lindqvist, and MW Fong. Modeling multistep cyber attacks for scenario recognition. In *DARPA Information Survivability Conference and Exposition*, pages 284–292, 2003.

Frédéric Cuppens and Alexandre Miège. Alert correlation in a cooperative intrusion detection framework. In *IEEE S&P*, 2002.

Oliver Dain and Robert K Cunningham. Fusing a Heterogeneous Alert Stream into Scenarios. In *ACM workshop on data mining for security applications*, pages 1–13, 2001.

Herve Debar and Andreas Wespi. Aggregation and Correlation of Intrusion-Detection Alerts. In *RAID*, pages 85–103, 2001.

S Zander and G Armitage. A survey of covert channels and countermeasures in computer network protocols. *Communications Surveys*, pages 44–57, 2007.

Yichi Zhang, Lingfeng Wang, Weiqing Sun, Robert C. Green II, and Mansoor Alam. Distributed Intrusion Detection System in a Multi-Layer Network Architecture of Smart Grids. *IEEE Transactions on Smart Grid*, 2(4):796–808, 2011.

Chenfeng Vincent Zhou, Christopher Leckie, and Shanika Karunasekera. A survey of coordinated attacks and collaborative intrusion detection. *Computers & Security*, 29(1):124–140, 2010.