DEPLOYMENT CONSIDERATIONS FOR INTRUSION DETECTION SYSTEMS IN ADVANCED METERING INFRASTRUCTURE

BY

DAVID RAYMOND GROCHOCKI JR

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Electrical and Computer Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2013

Urbana, Illinois

Adviser:

Professor William H. Sanders

ABSTRACT

Advanced Metering Infrastructures (AMIs) enable advanced bidirectional communication between utilities and smart meters deployed in the field, allowing consumption, outage, and price information to be shared efficiently and reliably. The addition of this new infrastructure, connected through mesh networks, has given rise to new opportunities for adversaries to interfere with communications and possibly compromise utilities' assets or steal customers' private information.

The goal of this thesis is to survey the various threats facing AMIs in order to identify and understand the requirements for a comprehensive intrusion detection solution. The threat analysis leads to an extensive set of failure scenarios that captures the attackers' key objectives and is used to extract the information required to effectively detect attacks. Using the information taken from the failure scenarios and knowledge of how encrypted communications can affect detection reliability, we explore possible intrusion detection system (IDS) infrastructures and discuss deployment considerations for each of them, paying particular attention to how well they can detect attacks. We also suggest that the widest coverage of monitoring for attacks can be provided by a hybrid sensing infrastructure that uses both a centralized intrusion detection system and embedded meter or dedicated standalone sensors.

To my family, for their love and support.

ACKNOWLEDGMENTS

I would like to thank my adviser, William H. Sanders, for his mentoring and support. I would also like to thank Robin Berthier for his advice and insight throughout my time as part of the PERFORM group. This work would not have been possible without their guidance. I would like to thank my fellow PERFORM members for their feedback and support, and for their friendship. I thank Jenny Applequist for her editorial assistance and support.

I am grateful to the researchers outside of the PERFORM group, here at the University of Illinois and at Fujitsu Laboratories, for working with me in accomplishing my goals. I would particularly like to thank Rakesh Bobba, Jun Ho Huh, Ahmed Fawaz, Edmond Rogers, Alvaro Cardenas, and Jorjeta Jetcheva. This work would not have been possible without them.

This material is based upon work supported by the Department of Energy under Award Number DE-OE0000097 as part of the Trustworthy Cyber Infrastructure for the Power Grid Center and was indirectly supported by Sandia National Laboratories. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.[1]

Finally, I thank my parents, friends, and family for the support and encouragement they have given me over the years and for having faith in me for all that I set out to do.

[1]This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

ADVISE	Adversary View Security Evaluation
AMI	Advanced Metering Infrastructure
DCU	Data Collection Unit
DDoS	Distributed Denial of Service
DIDS	Distributed Intrusion Detection System
FAN	Field-Area Network
HAN	Home-Area Network
IDS	Intrusion Detection System
LAN	Local-Area Network
MANET	Mobile Ad Hoc Network
NAN	Neighborhood-Area Network
PLC	Power Line Communication
WAN	Wide-Area Network
WLAN	Wireless Local-Area Network

CHAPTER 1

INTRODUCTION

The size and scope of the electrical grid in the United States are almost unfathomable. The U.S. Energy Information Administration estimates that more than 3,200 public utilities control the over 10,000 power-generating units that deliver electricity through tens of thousands of miles of transmission and distribution lines to millions of customers across the country [1]. The introduction of a new metering infrastructure to energy delivery systems is a significant change that requires a tremendous amount of planning. Laying the proper security foundation is important in this planning to ensure that the infrastructure will be resilient against the threat landscape targeting advanced metering infrastructure (AMI).

1.1 Motivation

Protecting power grid assets from computer attacks is a matter of national security and public safety. However, in some countries, including the United States, many portions of the power grid infrastructure are managed by private enterprises, which often operate under tight security budgets. It is essential to identify the requirements for a comprehensive monitoring solution that would enable utilities to gain situational awareness over the security state of their infrastructure. Utilities need to understand the risks of AMI deployments and the requirements for intrusion detection before they choose the monitoring architecture in which to invest.

According to a publicly-available FBI intelligence bulletin, there is evidence that some AMI deployments have already been hacked [2]. Smart meters are low-cost commodity devices that operate in locations lacking the physical security necessary to prevent tampering. While basic protective measures have been developed (e.g., tamper-evident seals), they may not

provide sufficient protection from attacks over the lifespan of the meters, which has been estimated to span several decades [3]. As the rate of AMI deployments increases, the number of possible security threats continues to grow, raising concerns from governments and privacy groups alike in regards to consumer privacy, grid reliability, and national security.

AMIs need an efficient and effective means of monitoring traffic in order to detect and respond to malicious activity on the network. Distributed intrusion detection systems (DIDSes) show promise, though most research on them focuses on optimizing the arrangement of nodes and communication among them. Few studies examine DIDSes in the context of smart grid technology, and fewer still try to identify the best method for distributing nodes and the challenge of monitoring encrypted traffic in the field, or so-called *neighborhood-area networks (NANs)*, in which smart meters are deployed.

1.2 Contribution

In this thesis, we present a collection of failure scenarios resulting from an extensive survey of the AMI threat landscape. We explore high-level attacker motivations as well as individual attack techniques to elicit specific information required for detection. A second contribution is an analysis of possible intrusion detection system (IDS) deployment options in AMIs, weighing factors such as coverage and detection reliability. Combined with the failure scenarios and required detection information, that analysis provides researchers and utilities with the knowledge they need to make informed decisions regarding the state of security in their AMI deployments. Another contribution is an analysis and discussion of key management considerations in AMIs. Many smart meters have the option of enabling encrypted communications, which presents a new set of problems in monitoring the network for attacks. We examine existing work and propose possible solutions with respect to the various deployment architectures.

This thesis is organized as follows. The remainder of Chapter 1 provides relevant background information on AMIs and explores related work. In Chapter 2, we present a collection of failure scenarios resulting from an extensive survey of AMI-specific threats and a detailed mapping to the information required for accurate attack detection. We describe four pos-

sible deployment schemes for IDSes in AMIs in Chapter 3 and discuss various deployment considerations for each of them. Chapter 4 presents key management considerations for encrypted AMI communications for each of the architectures described in Chapter 3. We conclude with Chapter 5.

1.3 AMI Overview

The role of AMI is to facilitate two-way communication between utility companies and metering devices in smart grid initiatives around the world. AMI networks are being seen more and more as a general-purpose communication infrastructure that can be used for a wide range of smart grid needs, including not only billing, but demand response programs and distribution automation applications. They benefit utilities and consumers alike by providing remote electricity usage readings (on-demand and periodic), electricity price information communication, alerts about outages and blackouts, the ability to remotely update meter firmware, and more. Some AMI systems even allow operators to remotely connect and disconnect customers from the electric grid. Some of these communications require reliable real-time delivery, while others have the option of being buffered and delayed without any negative consequences to the system. Since sensitive customer information is frequently sent over the grid, additional security and privacy requirements have emerged.

In order to support those requirements as well as a wide range of meter deployment topologies (e.g., from dense urban settings to sparse rural environments), meter manufacturers have developed highly flexible network architectures that can incorporate many different communication media. Typically, an AMI network consists of a multi-hop mesh network composed of several thousand smart meters, a smaller number of gateway devices, and other network devices that are connected to each other via wireless or power line communication (PLC) technologies. AMI networks enable communication among smart meters and devices attached to them as well as with servers in the utility company's network. The networks commonly adhere to the same network hierarchy (see Figure 1.1), with a wide-area network (WAN) connecting the utilities to a set of gateways in the field, and then NANs, also called *field-area networks (FANs)*, connecting gateways to meters. Meters themselves can be used

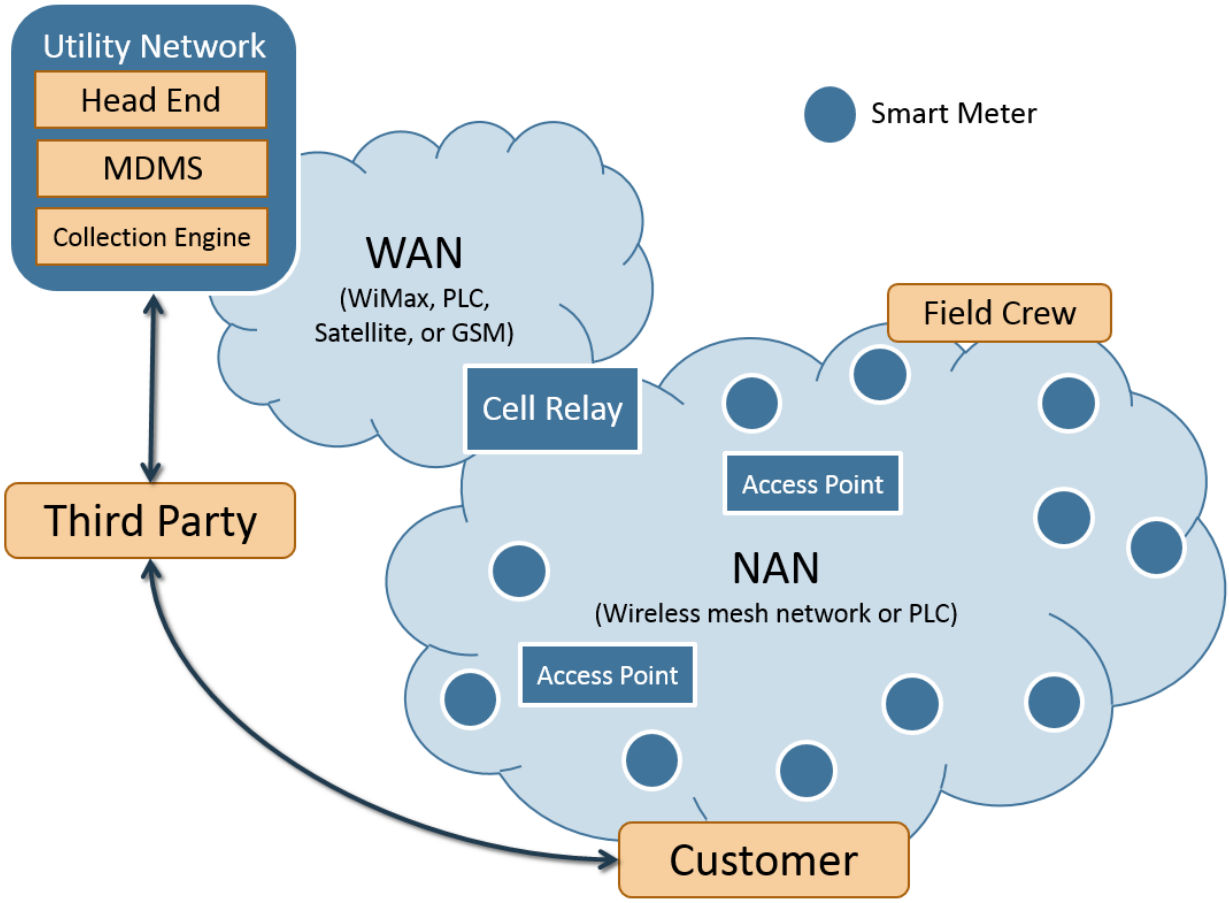

Figure 1.1: Typical AMI Architecture

as gateways to access the home-area network (HAN) deployed within customer premises to connect to thermostats and smart appliances.

A WAN uses long-range and high-bandwidth communication technologies, such as long-range wireless (e.g., WiMAX), cellular (e.g., 3G, EVDO, EDGE, GPRS, or CDMA), satellite, or PLC. NANs typically have shorter range requirements and can be deployed using wireless (e.g., IEEE 802.11, IEEE 802.15, or proprietary communication stacks) or PLC-based technologies. In some cases, meters can directly include cellular capabilities or even use the customer's home Internet connection to bypass the need for separate WANs and local-area networks (LANs). In this thesis, we focus on NANs that use a wireless mesh network. The mesh topology brings robustness to the network, since communication routes can automatically adapt when failures occur. However, they also represent a challenge for the deployment

of an efficient security monitoring solution, because of their distributed nature and their use of wireless communication technologies.

1.4 Related Work

The continuously growing threat landscape for AMIs has grabbed the attention of many security researchers, and some of their recent work is closely related to ours. We conducted a thorough survey of previous literature to identify fundamental attack steps within AMIs and HANs.

1.4.1 Threat Analysis in AMI

A first category of research we studied did not specifically cover AMIs, but was useful in clarifying the threats common to wireless networks. For instance, [4] examines attacks on wireless networks to motivate solutions to address the privacy issue; [5] develops a threat model to guide the design of a secure wireless local-area network (WLAN) architecture; [6] and [7] study threats on mobile ad hoc networks (MANETs); [8] focuses on sensor networks; and [9] and [10] investigate threats specific to mesh networks.

In the category of publications focusing on the smart grid, [11] presents the design of a firewall to secure wireless communication in energy delivery systems. [12] examines attacks targeting energy theft in AMIs; the authors later used that analysis to motivate a new methodology for penetration testing in AMIs [13]. [14] offers a more comprehensive survey of the threat landscape for AMIs. Some research efforts have addressed the security of the wireless protocols popular in HAN communication. For instance, [15] surveys popular security protocols used in wireless sensor networks for authentication, including ZigBee, SPINS [16], and LEAP [17]. In-depth security analyses for these protocols, like ZigBee in [18] and [19] and LEAP+ in [20], have been performed as well.

Home automation is an interesting area of research that utilizes many of the previously mentioned wireless protocols. [21] explores wireless home automation, while discussing interoperability between local home automation and the Internet. The use of a multi-port

power electronic interface (MPEI) for smart home automation is discussed in [22], which examines the systems' ability to interact with the home (e.g., monitoring and controlling home sensors or controlling appliances) through the use of an ethernet-connected interface (such as a smartphone or laptop). [23] explores the EPIC open source home automation framework with regard to interoperability, reliability, and support for multiple protocols, including ZigBee. However, while those publications address usage and functionality of home automation systems, they largely ignore the security aspects, even though they address the benefits of such systems being connected to the Internet.

Still, there are some researchers addressing the security concerns of HANs. Customer privacy issues, including the collection of energy consumption data to construct user behavior patterns, are discussed in [24], [25], [26], and [27]. Enhancing security in HANs has been explored in [28], which proposes a new authentication mechanism for smart devices. [29] identifies a set of IDS requirements for AMIs and briefly mentions different sensor deployment locations, including dedicated and meter-level sensors. [30] expands on that work to develop a specification-based IDS specifically targeting attacks in HANs. Other closely related work is described in [14], which covers threats prevalent in the AMI mesh network, along with several sensor deployment ideas, suggesting that a hybrid approach (such as we describe in Section 3.5) would provide the widest monitoring coverage.

CHAPTER 2

FAILURE SCENARIOS

As the new communication and computational capabilities of smart grid devices are being added to traditional energy delivery systems, the attack surface is increasing significantly. Attacks on utilities that would normally require physical access to the utility network from a more centralized location may now be possible through a remote exploit executed over the Internet via access to one of the new devices. Smart meters are not just connected to utility networks, but also directly connected to the devices and networks in customers' homes, which, in turn, are connected to the Internet and cellular networks. As those networks, which were once isolated from each other, are coming together, new attack possibilities may appear. In the context of AMIs and HANs, meters appear to be an attractive target for adversaries, because not only are they deployed in large numbers and lacking the necessary physical protection, but also they typically have limited or no security monitoring capabilities.

The goals of this chapter are to review the possible attack motivations and failure scenarios that are specific to AMI networks and to tie them to individual attack techniques. A number of representative case studies are explored to connect the attackers' motivations with more fine-grained, individual attack techniques. The results lead to an extensive collection of failure scenarios and to the identification of the information required for detecting such scenarios. Note that our analysis focused on AMI networks. While we include HAN access as a possible entry vector, we do not explore HAN vulnerabilities in depth, and there may be additional attacker motivations that involve additional attack techniques related to HANs; however, these topics are beyond the scope of this thesis.

2.1 Threat Survey

The key characteristics of an AMI that could attract malicious activity are (1) access to a communication infrastructure other than the Internet, (2) access to millions of low-computation devices, (3) access to sensitive customer information, (4) high visibility and high impact in the case of disruption (e.g., a power outage), and (5) the financial value of energy consumption data. Consequently, attackers could be motivated to abuse the communication infrastructure, reduce their energy bills, steal information from targeted customers, remotely disconnect targeted customers or large regions, or create denial-of-critical-services.

A large set of attack techniques can be combined to reach those objectives. We conducted a thorough survey of previous literature from 11 different universities and corporations to identify fundamental attack techniques within AMI networks.

Our next step was to combine the attacks discussed in the literature in order to build a holistic view of the AMI failure possibilities. From an initial list of 5 attack motivations and more than 30 unique attack techniques, we first filtered out those irrelevant to the AMI environment, and then worked on decomposing the remaining ones into individual attack components. The motivation for the decomposition was to understand the fundamental pieces of information required by a monitoring solution to detect any combination of those attack techniques, including combinations that we did not cover in our threat model. We accomplished this in [14] for failure scenarios in AMI networks; however, HANs were outside the scope of our initial analysis.

We expand on that previous survey here to include scenarios and motivations affecting not only AMI networks but also HANs. Our list of attack techniques grew, and we illustrate the decomposition through a series of case studies. Finally, we present an updated set of failure scenarios by combining the original set of scenarios presented in [14] with the additional techniques we learned from our analysis of HANs.

2.2 Attacker Motivation

We will now explore each of the high-level motivations of the adversaries in the failure scenarios to follow:

Destroy Equipment

Destroying a utility's power generation, transmission, or distribution equipment has high public visibility and can be both costly to utility companies and inconvenient to consumers. Whether the damage is to large transformers, neighborhood substations, or individual smart meters or appliances, the end result can be long repair or replacement times and loss of public confidence in the smart grid, in addition to financial strain on the companies managing the grid.

Interrupt Service

While interrupting service can be achieved by destroying key equipment, there are alternative methods of causing blackouts and other power interruptions that do not involve physical destruction of electrical equipment. Service interruptions can affect not just power, but also the financial and billing systems of the utility, and can cause load-balancing problems in the management of electrical power demands.

Information Gathering

As advanced communication abilities are being added to electrical distribution in smart grid deployments, much more information is being sent between customers and utilities. Customer information being transmitted back to the utility poses a privacy concern and can be used maliciously by someone eavesdropping on the network. Additionally, as billing and usage information are communicated over the same network, it may be possible for an attacker to access the utility's network from any point in the AMI network. Finally, if any key distribution system is in place, monitoring of traffic may provide adversaries with encryption keys, if proper security mechanisms are not in place.

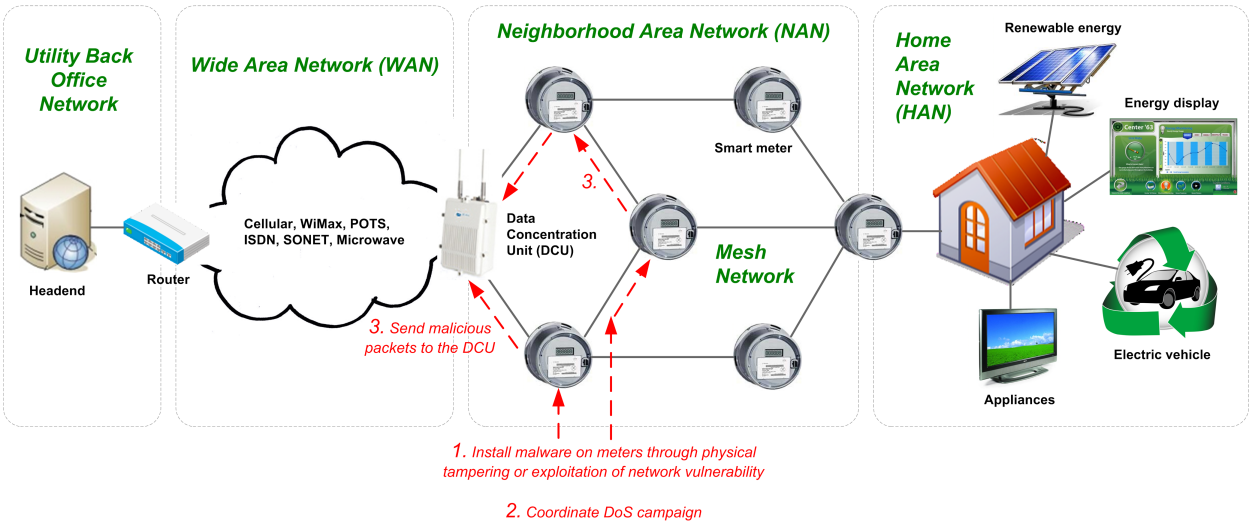

Figure 2.1: Distributed DDoS Attack on a DCU [14]

Communication Medium

Whether it is used to coordinate additional attacks or to serve as a covert communication channel, the infrastructure of the AMI network may be useful to adversaries.

Energy Theft

Energy theft is a big concern for utility companies. It may consist of a small-scale operation to lower a home electric bill or a larger effort to financially damage the utility more severely. Now that everything in the smart grid is digital, there are new opportunities for adversaries to control the system.

2.3 Case Studies

To elicit information that is needed to detect individual attacks, we explored a comprehensive set of case studies that are representative of the attackers' key objectives. Rather than include every case study, we present a representative subset of them in this section and include the results of each of them in the failure scenarios that follow. In the remainder of this section, we look at five example case studies to demonstrate how they were used to tie the attackers' key objectives to the attack techniques.

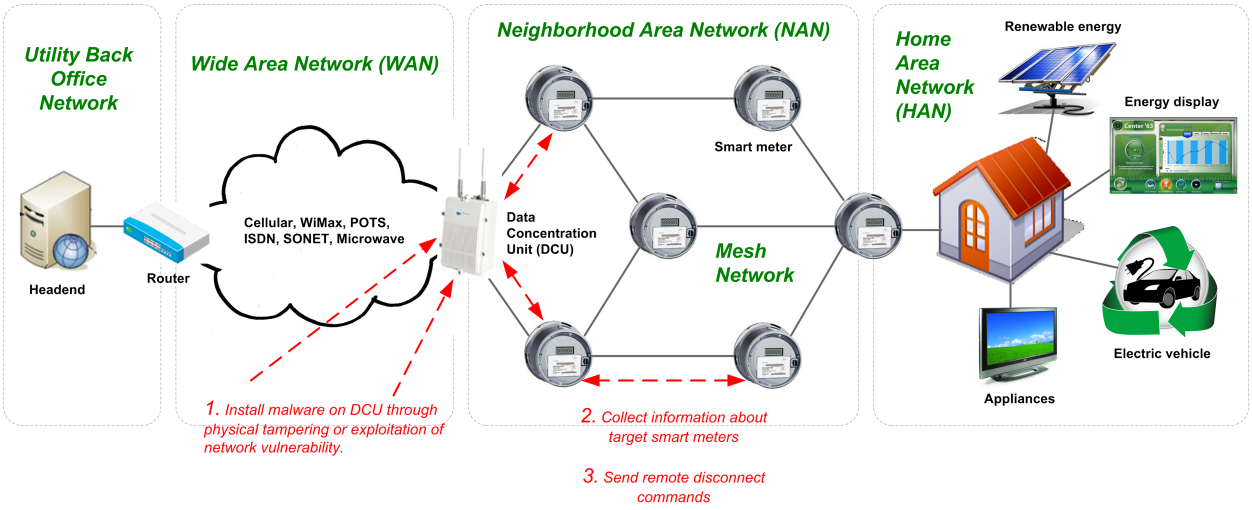

Figure 2.2: Sending of Remote Disconnect Commands through the DCU [14]

DDoS Attack Against the Data Collection Unit (DCU)

The attacker's motivation in performing a distributed denial of service (DDoS) attack [31] is to compromise the DCU and prevent relays deployed between the WAN and NANs from communicating or functioning (see Figure 2.1). Assuming that the entry point of the attack is in the smart meters, the following are typical individual attack techniques that would be involved: (1) installation of malware on the meters through physical tampering or exploitation of a network vulnerability; (2) coordination of a DDoS campaign among the compromised meters; and (3) crafting and sending of a large number of malicious network packets to the DCU.

Stealing Customer Information

The motivation of the second attack is to collect customer information and learn about customer behavior by eavesdropping on the incoming and outgoing network traffic of the meters. Considering that the AMI traffic may be encrypted, this attack may involve the following individual steps: (1) theft of the decryption keys (or the master seed number that is used to generate the keys), accomplished by physically accessing the meters or performing brute-force attacks on the crypto system; (2) eavesdropping on the AMI traffic to intercept the messages; and (3) decryption of the messages and collection of the message contents.

Sending Remote Disconnect Commands Through the DCU

Here, the attacker wishes to disconnect a large number of customers by exploiting the "remote disconnect" functionality on the meters (see Figure 2.2). The DCU is very likely to be the point for launches of these attacks, as it is one of the more suitable devices for triggering the remote disconnect command for many customers without being detected by the utility. The attack techniques involved are (1) installation of malware on the DCU through physical tampering, exploitation of a network vulnerability, or abuse of insider privileges; (2) identification of the meters and collection of information about them (e.g., IP addresses); and (3) sending of remote disconnect commands to the targeted meters.

Remote Access to Smart Devices

The motivation for this attack is to gain remote access to smart devices (e.g., home appliances, electric vehicles, energy dashboards), perhaps to disable appliances (e.g., turn off refrigeration units in a warehouse, resulting in the loss of products). Assuming that the entry point of the attack is a compromised smart meter (see Figure 2.3), the following are typical individual attack techniques that would be involved: (1) installation of malware on the smart device through the exploitation of a network vulnerability; and (2) sending of a power-off or shut-down command to the smart device.

Attacking a NAN from a HAN

In this attack, the motivation for compromising the NAN from the HAN is to yield additional attack vectors, if compromising a smart meter directly is not an available option (see Figure 2.4). Assuming that the entry point of the attack is through the Internet, the following are typical individual attack techniques that would be involved: (1) installation of malware on a smart device through the exploitation of a network vulnerability; (2) installation of malware on a meter through the exploitation of a network vulnerability; and (3) communication through the smart device to carry out attacks from the smart meter.

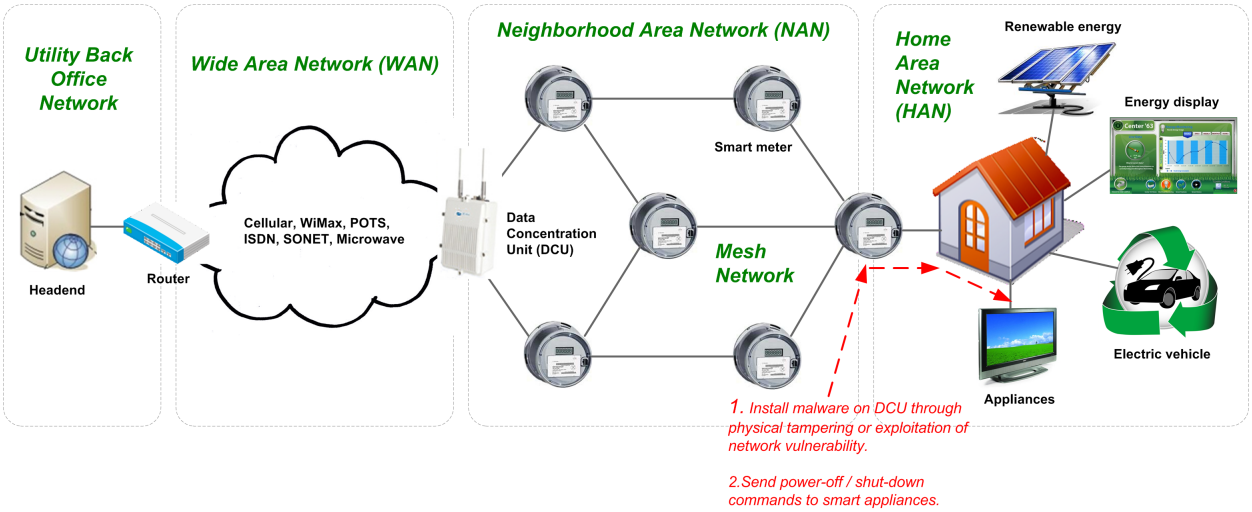

Figure 2.3: Attacking a HAN from a NAN

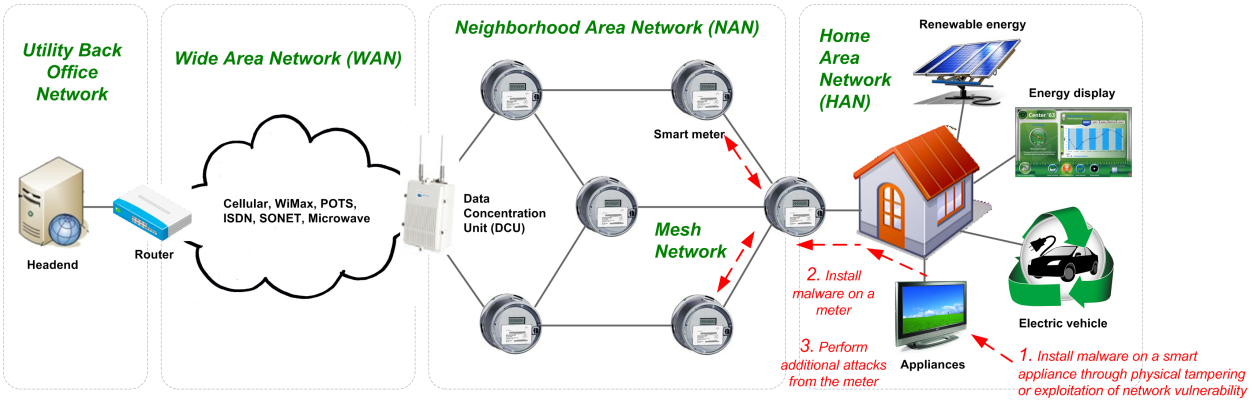

Figure 2.4: Attacking a NAN from a HAN

2.4 Failure Scenarios

Based on the results collected from the previously discussed threat survey and case studies, we created a collection of failure scenarios that encompass the threat landscape of AMIs. Next, we describe those scenarios, organized according to the attacker motivations detailed in Section 2.2. We begin by detailing intermediate goals that adversaries would need to attain before continuing to the main objectives.

2.4.1 Failure Scenario Template

The graphical notation used is illustrated in Figure 2.5 and shows the various representations that are used throughout the set of failure scenarios. Green nodes are the entry points to the scenario. Green rectangles represent various ways an adversary may gain access to the network (e.g., physical access to a meter through the optical port), while green ovals signify a common intermediate goal that can be achieved through some other set of attack techniques (e.g., access to a single compromised meter via a remotely executed network exploit).

Every failure scenario is terminated by an orange oval, which represents the attacker's motivation or main objective, or a green oval, which represents an intermediate goal whose achievement enables additional actions. Intermediate goals can reappear in other scenarios. The blue parallelograms represent individual attack techniques. Connection by a dotted line indicates an "OR" relationship; that is, the child step can happen if any of the parent steps occur. Solid lines indicate an "AND" relationship, which indicates that the child step can follow only if all of the parent steps occur.

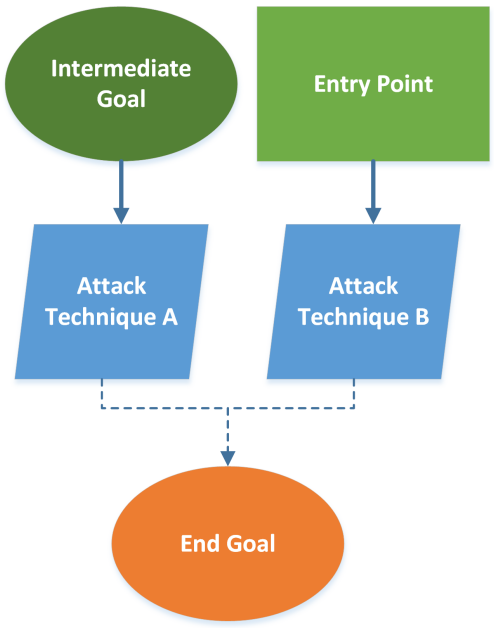

Figure 2.5: Failure Scenario: Template

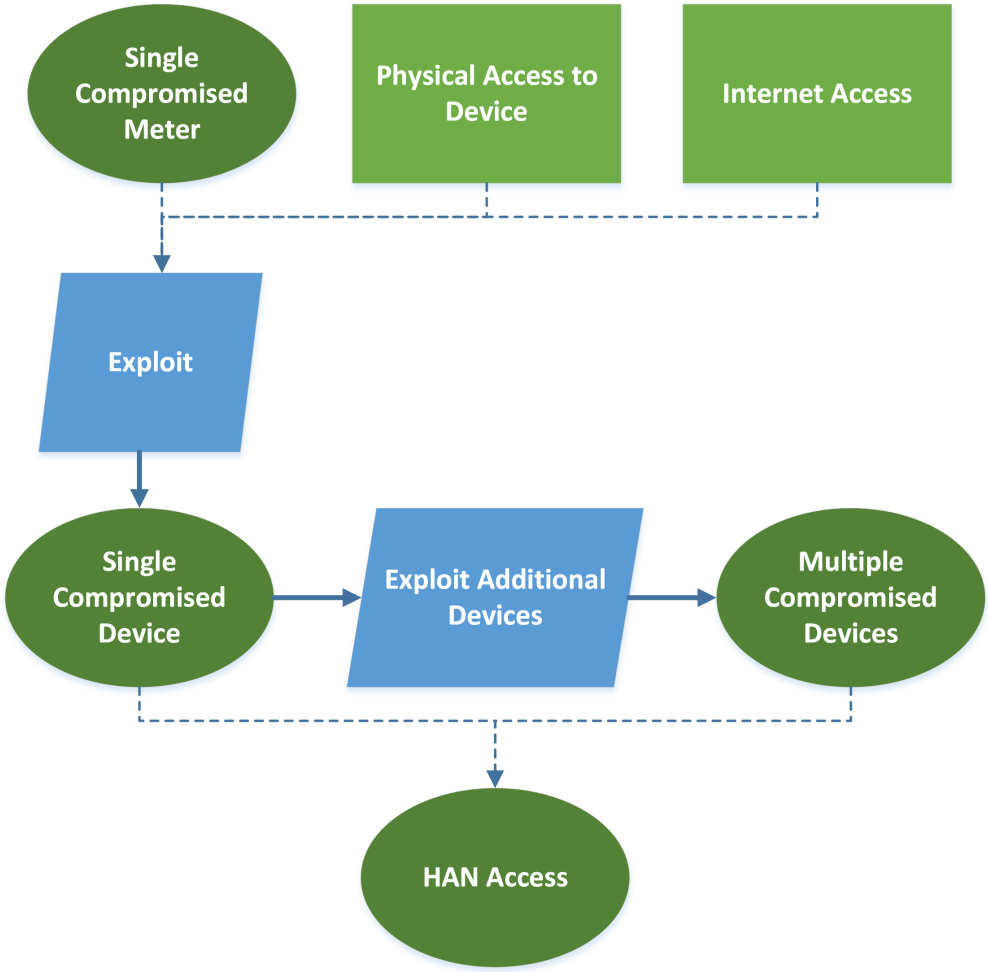

Figure 2.6: Failure Scenario: Gaining Access to a HAN

2.4.2 Building Blocks

Here, we present a set of failure scenarios that do not result in achievement of any of the possible attacker goals discussed previously, but rather show some of the important intermediate goals utilized in the remainder of the scenarios. Figure 2.6 shows the possible set of entry points and attack techniques that can be used to gain the intermediate goal of accessing a HAN, while Figure 2.7 does the same for achieving the intermediate goal of compromising one or more meters, and Figure 2.8 for eavesdropping. To understand how one may go about constructing an attack utilizing these scenarios, take, for example, the intermediate goal of "Multiple Compromised Meters" in Figure 2.7. It is possible to construct one possible attack by first gaining physical access to a meter and then using

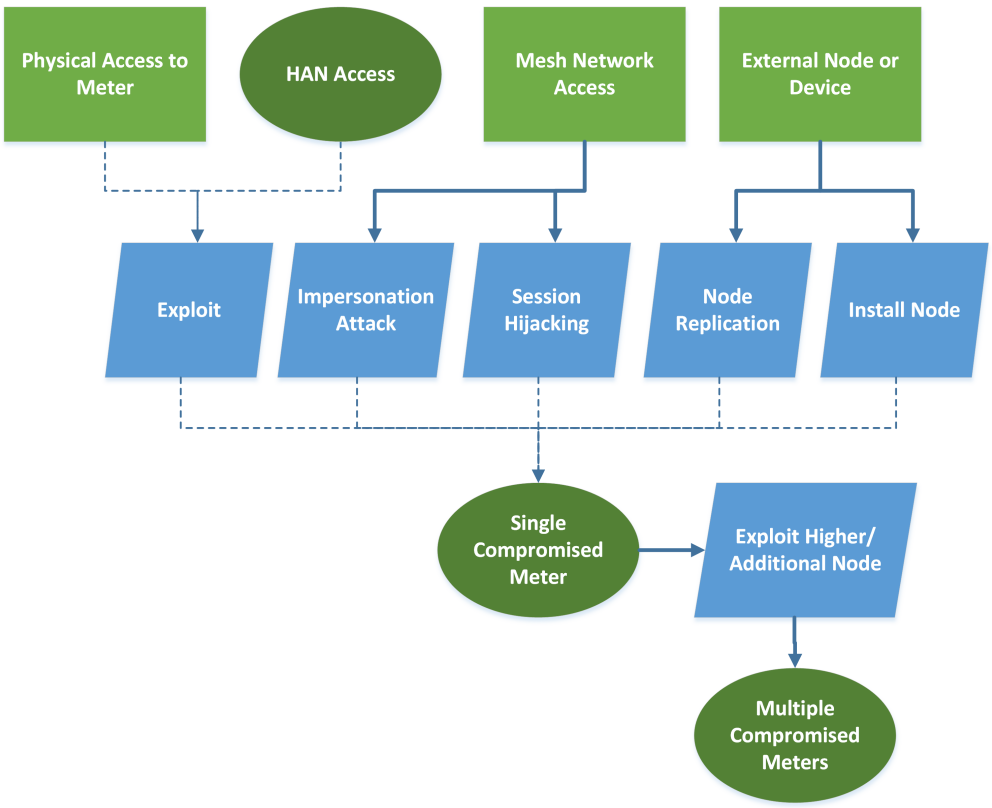

Figure 2.7: Failure Scenario: Compromising One or More Meters

a buffer overflow attack (exploit) to gain root access, thus compromising a single node. Once a single meter has been compromised, a network exploit could be used to compromise multiple meters, thus achieving the desired intermediate goal. However, one could have used a different entry point and set of attack techniques to accomplish the same goal.

Figures 2.6 and 2.7 also utilize all of the unique entry points (represented by green rectangles) used throughout the failure scenarios:

Physical Access to Meter

Smart meters operate in locations lacking the physical security necessary to prevent tampering. While basic protective measures have been developed (e.g., tamper-evident seals), they may not provide sufficient protection from adversaries looking to gain entry to the network, whether by circumventing any tamper-resistant protections, interfacing with the optical port on the meter, or communicating wirelessly with the device; physical access to a meter may be the most accessible point of entry for any adversary.

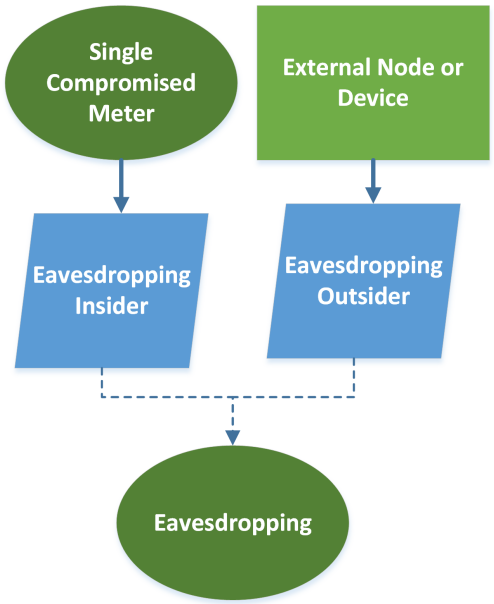

Figure 2.8: Failure Scenario: Eavesdropping

Physical Access to Device

Smart devices with network access, like thermostats, refrigerators, and smart cars, are becoming increasingly popular in homes across the country. As these devices begin to communicate with smart meters, they open up new access points to the network from within the home. While these devices can be protected better than smart meters by being placed indoors, they still serve as potential points of entry into the AMI network.

Internet Access

Internet access is easy to attain and can provide some degree of anonymity. Unlike direct access to a smart meter or smart device, however, Internet access requires that the adversary have additional knowledge in order to breach the boundary of the AMI network. Unless the AMI utilizes the Internet in WAN communications, the Internet might be a useful entry point in targeting AMI through the HAN.

Mesh Network Access

Accessing the mesh network that makes up the AMI communication is more difficult than simply connecting to the Internet. The attacker may require specialized communication equipment and wireless radios in order to communicate with devices in the

AMI. That could be costly, and it would require in-depth knowledge of the wireless technologies and protocols used in communication over the network.

External Node or Device

Control of a custom meter or network node might be difficult to achieve, but may prove to be a useful asset in gaining access to the AMI network. Other devices, like wireless signal jammers, might be useful in achieving objectives as well.

2.4.3 Destroy Equipment

Figure 2.9 shows the entry points and attack techniques that can be used by an attacker with the *Destroy Equipment* motivation. Damage to equipment can include destruction of transformers, overloading of power substations, bricking of individual smart meters, and communication failures, among other things. In turn, they can lead to high time and monetary costs to the utility and can be highly visible to the public, resulting in long-term lack of confidence in the smart grid and loss of reputation to the utility. This failure scenario also introduces the *Alter Network Stream* intermediate goal, which involves techniques that change the information that is being communicated over the AMI network. It is important to note that there may be additional kinetic methods of achieving the same goal, but those methods are outside the scope of this thesis.

Entry Points

- Single compromised meter (more than one meter may be involved, but a minimum of one is required)
- Mesh network access

Impact

- Repair or replacement cost to utility
- Financial damage due to inability to bill customers
- Potential for brownouts or blackouts

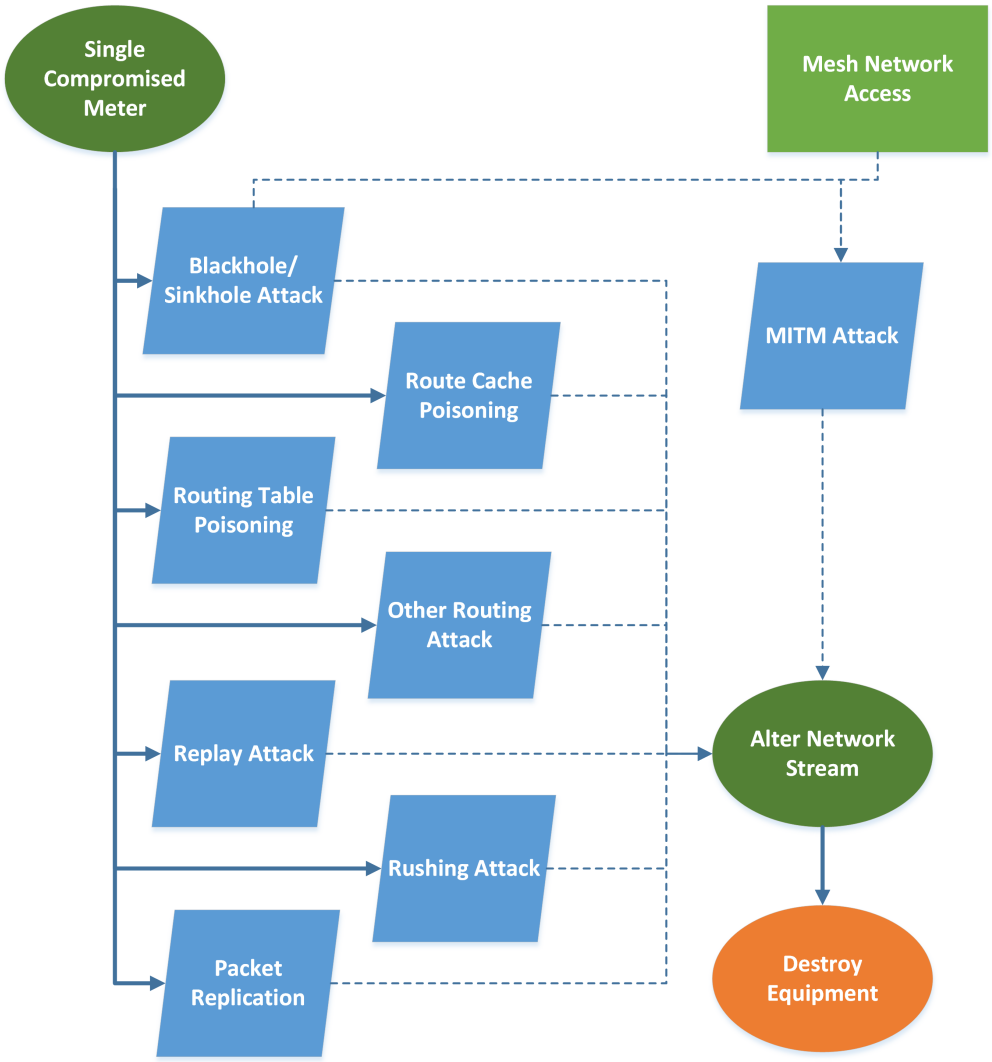

Figure 2.9: Failure Scenario: Destroy Equipment

- Potential cascading failures resulting in larger crisis
- Possible loss of life
- Public relations problems due to damaged reputation
- Loss of public confidence in utility and smart grid

2.4.4 Interrupt Service

Figure 2.10 shows the entry points and attack techniques that can be used by an attacker with the *Interrupt Service* motivation. Adversaries may find ways to interrupt service beyond

causing power outages, including damaging a utility's ability to respond to cyber events. Bringing down communication infrastructure has impacts on billing and power generation and distribution systems. Interruption of service has an impact similar to that of destroying equipment, but can be accomplished more covertly, if so desired.

Entry Points

- Single compromised meter
- Multiple compromised meters
- Physical access to meter
- External node or device

Impact

- Repair or replacement cost to utility
- Financial damage due to inability to bill customers
- Loss of ability to quickly respond to threats
- Loss of ability to measure and respond to changes in demand
- Potential for brownouts or blackouts
- Potential cascading failures resulting in larger crisis
- Public relations issues due to damaged reputation
- Loss of public confidence in utility and smart grid

2.4.5 Information Gathering

Figure 2.11 shows the entry points and attack techniques that can be used by an attacker with the *Information Gathering* motivation. Information gathering can be used to gain access to encryption keys, private utility information, or personal customer information, among other things. Covert compromise of the key distribution system of a utility can open up a large array of other attack vectors and lead to other potential problems. Data breaches

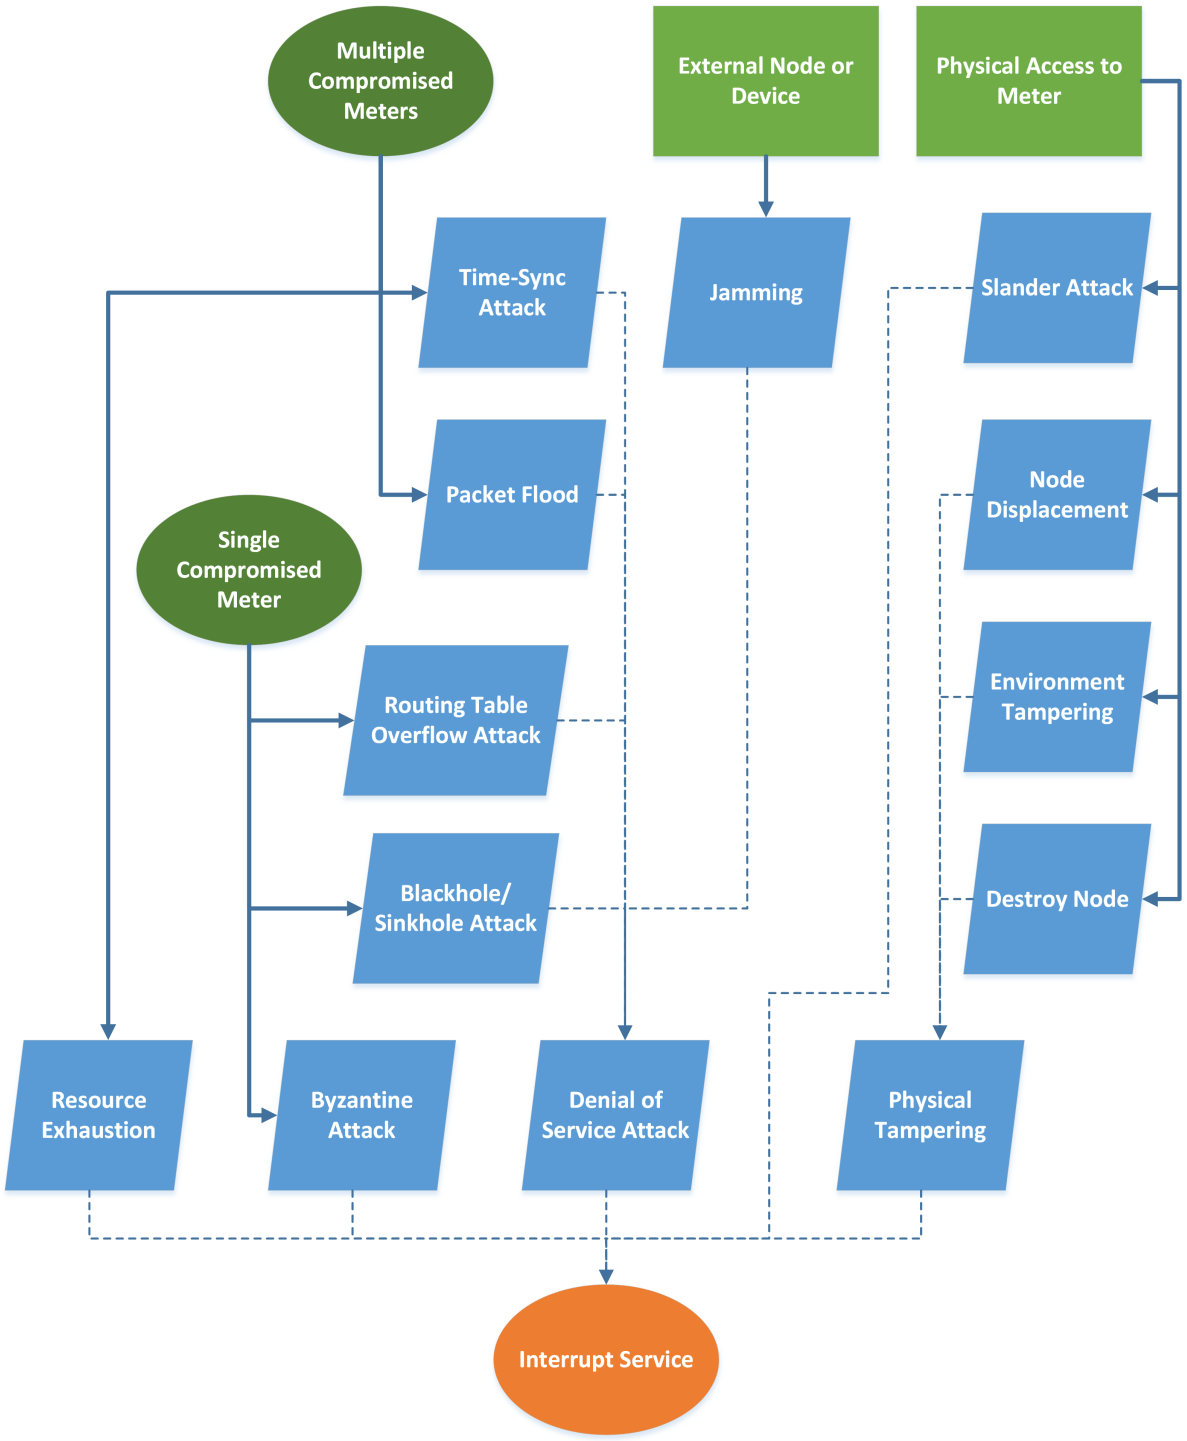

Figure 2.10: Failure Scenario: Interrupt Service

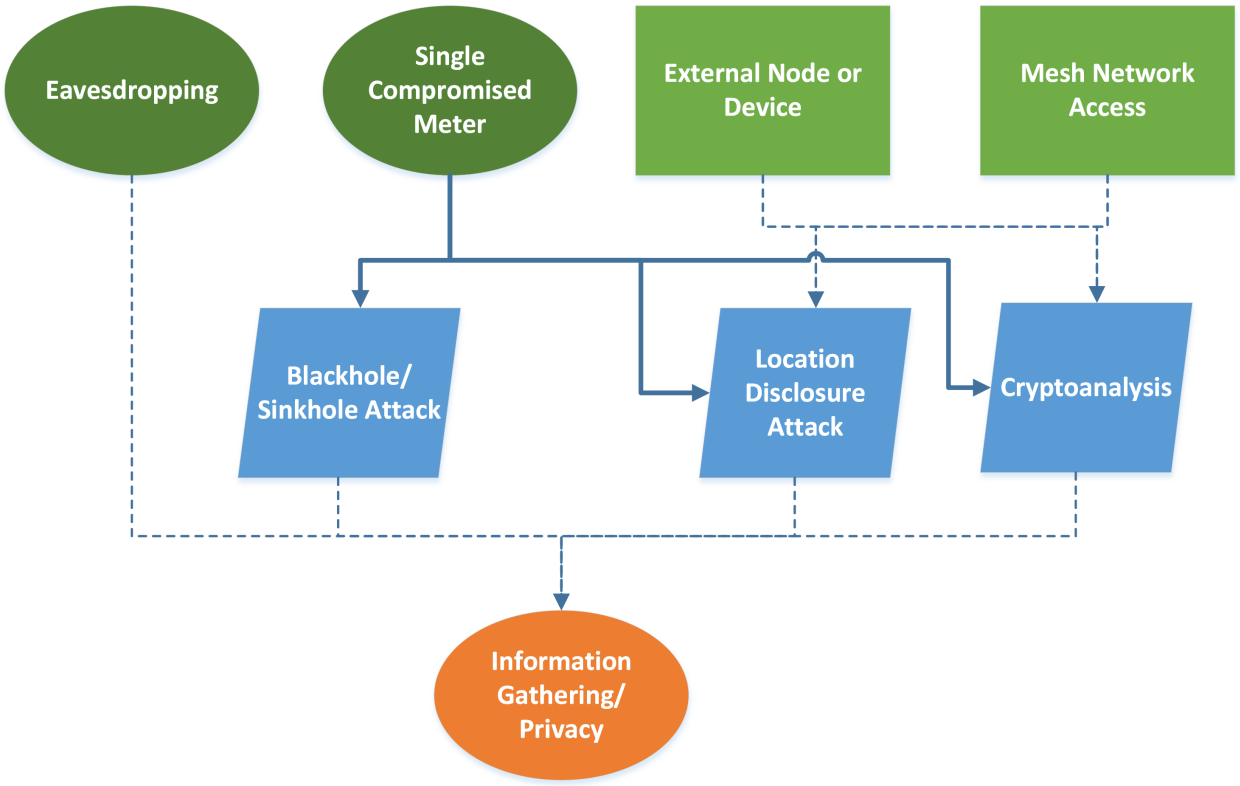

Figure 2.11: Failure Scenario: Information Gathering

can result in high costs to the utility if customer privacy is compromised and can lead to lack of public confidence in the smart grid and damage to the reputation of the utility.

Entry Points

- Single compromised meter (more than one meter may be involved, but a minimum of one is required)
- Mesh network access
- External node or device
- Eavesdropping

Impact

- Key compromise may require costly rekeying effort
- Financial damage due to privacy breach

- Financial damage due to lawsuits and fines
- Customers may experience privacy violations
- Public relations issues due to damaged reputation
- Loss of public confidence in utility and smart grid

2.4.6 Communication Medium

Figure 2.12 shows the entry points and attack techniques that can be used by attackers with the *Communication Medium* motivation. While using the AMI network strictly as a covert communication medium may not directly harm the utility, it may degrade the utility's situational awareness of the network by causing delays in regular communications. It can also be used to coordinate additional attacks that may cause more direct harm to the utility financially or to the quality of service delivered to customers.

Entry Points

- Multiple compromised meters

Impact

- Loss of ability to quickly respond to threats
- Loss of ability to measure and respond to changes in demand

2.4.7 Energy Theft

Figure 2.13 shows the entry points and attack techniques that can be used by an attacker with the *Energy Theft* motivation. Energy theft can result in large financial losses to the utility due to lack of revenue. It can be accomplished on a small scale that slowly affects the financial health of the utility or on a larger scale with high visibility, which may result in loss of public confidence in the smart grid. Alternatively, adversaries may reverse the outcome by providing the utility company with higher revenues. That could result in financial loss to

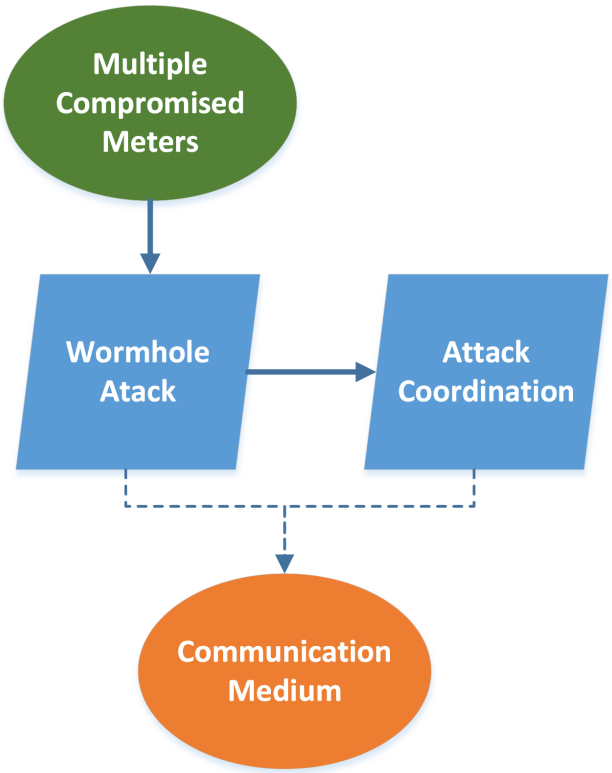

Figure 2.12: Failure Scenario: Communication Medium

the utility in terms of lawsuits and fines, and could be even more damaging to the utility's reputation.

Entry Points

- External node or device
- Eavesdropping
- Alter network stream

Impact

- Financial damage due to loss of revenue
- Financial damage due to lawsuits and fines if the utility is found to be overcharging
- Potential for brownouts or blackouts
- Customer inconvenience

- Customer service issues due to actions required to resolve billing disputes
- Public relations issues due to damaged reputation
- Loss of public confidence in utility and smart grid

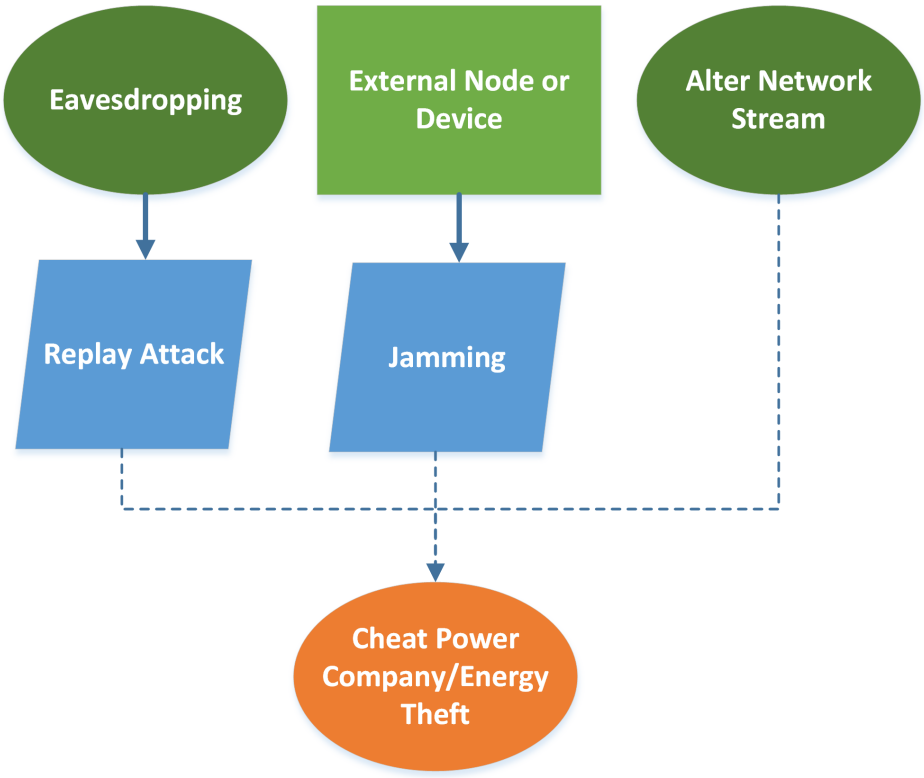

Figure 2.13: Failure Scenario: Energy Theft

2.5 Information Required for Detection

From the failure scenarios presented in Section 2.4, we were able to extract individual attack techniques. The set of information required to detect different steps in those attack techniques is presented in Table 2.1.

Each row of the table is an individual low-level attack technique that can be used alone or in combination with other techniques to build complex failure scenarios. As explained in the previous section, our goals are to identify fundamental attack techniques and to identify the core information required for their detection. With that information, we will be able to

Table 2.1: List of individual attack techniques and information required to detect them (adapted from [14])

Category	*Attack technique*	*Target*	*Information required*
DoS	Collision in Packet Transmission	NAN Link Layer	NAN collision rate, node response time
DoS	Packet Flood	Node in NAN (Meter/DCU)	CPU and memory usage of target incoming network traffic, authorized network protocols, network health information, packet-per-second rate, node response time
DoS	Jamming	NAN Physical Layer	NAN signal level, node response time
DoS	Alter Routing Table	Routing Protocol	Routing table health, node response time
DoS	Drop Packets	NAN Traffic	Packet loss among nodes in mesh network
DoS	Destroy Node	Node in NAN (Meter/DCU)	Node availability / response time
DoS	Time-Desynchronization	Node in NAN (DCU)	Time-synchronization traffic among nodes or time configured on nodes
DoS	Resource Exhaustion (Battery, Bandwidth, or CPU)	Node in NAN (Meter/DCU)	Traffic among meters, valid traffic profile or node health (CPU, battery consumption), network health (bandwith usage)
Spoofing	Impersonate Regular Node	Node in NAN (Meter)	Associations between physical addresses and node identity
Spoofing	Impersonate Master Node	Node in NAN (DCU)	Associations between physical addresses and node identity, associations between regular and master node registrations
Spoofing	Man-in-the-Middle	NAN Traffic	Associations between physical addresses and node identity
Spoofing	Wormhole	NAN Traffic	Associations between physical addresses and node identity, routing table integrity/update
Spoofing	Slander	Distributed Detection System	Integrity of trust and reputation system
Eavesdropping	Passively Listen to Traffic	NAN Traffic	N/A (undetectable)
Eavesdropping	Active Cryptoanalysis	NAN Traffic	Traffic among meters
Physical	Compromise Meter	Node in NAN (Meter)	Integrity of meter firmware, memory contents of meter, meter firmware upgrade policy, meter status, information about bandwidth and wireless signal
Communication	Attack Coordination	Traffic in NAN	Network protocols that are authorized for use, network traffic among the meters, network characteristics of legitimate traffic

ensure that any combination of those steps can be detected when the appropriate detection technology is used.

The information required for detection can be organized into three categories:

System Information

Health reports from meter and gateways (CPU, battery consumption), firmware and software integrity of AMI devices, clock synchronization

Network Information

NAN collision rate, packet loss, node response time, traffic rate, health and integrity of routing table, associations between physical addresses and node identity

Policy Information

Authorized AMI protocols, authorized AMI devices, authorized traffic patterns, authorized route updates, authorized firmware updates

It is crucial to extract knowledge from the mapping between attacks and information required for detection in order to design a comprehensive and cost-efficient monitoring solution. Certainly, the categorizations presented in the table reveal that data must be collected from multiple points in the infrastructure. For example, the need for information on the health and integrity of routing tables requires routers (or, in the context of AMIs, meters) to be instrumented so that they can send periodic health reports or at least be remotely queried for health and integrity checks. However, instrumentation of all routers in the network may be too expensive, and a more cost-effective solution could be to rely on attack manifestations at other locations in the system instead of routers for detection.

CHAPTER 3

IDS ARCHITECTURES

Having identified the information required to detect common attacks in the previous chapter, we are now in a position to sketch possible IDS deployment schemes. This chapter explores four different approaches and discusses how effective each would be in detecting attacks, as well as deployment considerations for each of them.

3.1 Architecture Overview

The four architectures discussed in this chapter all have benefits and drawbacks that must be weighed before the most appropriate security monitoring solution is selected. We present an overview of each of the architectures before discussing deployment considerations for each.

Centralized Infrastructure

This option is both cost-effective and well-tested, as it is similar to many enterprise network security monitoring solutions. It provides many benefits in terms of cost and visibility, especially for encrypted traffic, but fails to provide full situational awareness over the entire AMI network.

Embedded Infrastructure

An embedded infrastructure provides the most comprehensive coverage of meter-to-meter communication and monitoring of smart meter devices. However, key management concerns and computational limitations are drawbacks of this type of deployment.

Dedicated Infrastructure

This solution provides wide coverage over most of the AMI network and is better able to handle key management concerns. Despite those benefits, this type of infrastructure

has added complexities due to placement of the sensor units, and it is not ideal for detecting attacks that aim to compromise any single device.

Hybrid Infrastructure

By combining the centralized and embedded or dedicated infrastructures, we are able to reap the benefits from multiple solutions. However, this approach increases the overall cost of the IDS deployment, and it does not remove all of the limitations of the various architectures.

3.2 Centralized Infrastructure

A centralized monitoring architecture is the most cost-effective solution and would consist of a single IDS sensor deployed at the head-end located in the utility data center, as shown in Figure 3.1. That is where smart meter data is processed and stored, and the sensor would be able to monitor all traffic that flows in and out of the AMI through the data center. In such a deployment, this type of infrastructure would only be able to access and analyze network traffic to and from the AMI network, maintenance and upgrade policies, and system logs from AMI appliances; peer-to-peer traffic between nodes in the AMI network would not be visible. Therefore, this type of sensor would be able to detect systemic attacks that target the utility network and insider attacks that leave traces in access logs. It would also be able to analyze anti-tampering alerts sent to the utility by smart meters.

The centralized infrastructure approach is similar to many enterprise solutions employed today on corporate networks, where data is transmitted between two different networks (e.g., Intranet-Intranet or Intranet-Internet). In such setups, data can be monitored at a central location, since communication typically has to go through a small number of routers, or endpoints, that serve as points of traffic aggregation in the network. One advantage of this approach is that the utility would be able to perform analysis on fully decrypted packets, as an IDS sensor located at the data center would have access to all the cryptographic keys needed to decrypt encrypted traffic.

Even though a centralized IDS infrastructure would be able to capture most of the traffic

in the network, its detection capabilities would be limited to detecting attacks that originate from the utility, targeting the AMI network, or vice versa. While our analysis of the information required for detection (see previous chapter) shows that an IDS sensor at the head-end is needed, it is not sufficient to provide an encompassing situational awareness of the AMI network as a whole. Certainly, there could be a number of attacks within the AMI that do not pass through the utility data center.

Attacks against the routing protocol of the mesh network, MAC or PHY layer attacks, and end-to-end application layer attacks between peer-to-peer AMI nodes are examples of items that a central sensor would miss. For instance, attack techniques such as "installing malware on the meter" or "eavesdropping on NAN traffic" through an active cryptanalysis (i.e., by injecting traffic to force nodes to generate encrypted packets) would be undetected by the central sensor, because it would not have access to information such as the integrity of the meter firmware, the memory contents of the meter, the NAN traffic among meters, the network bandwidth usage, or the routing table integrity. Such gaps in attack coverage associated with the centralized IDS warrant a closer look at a more distributed approach.

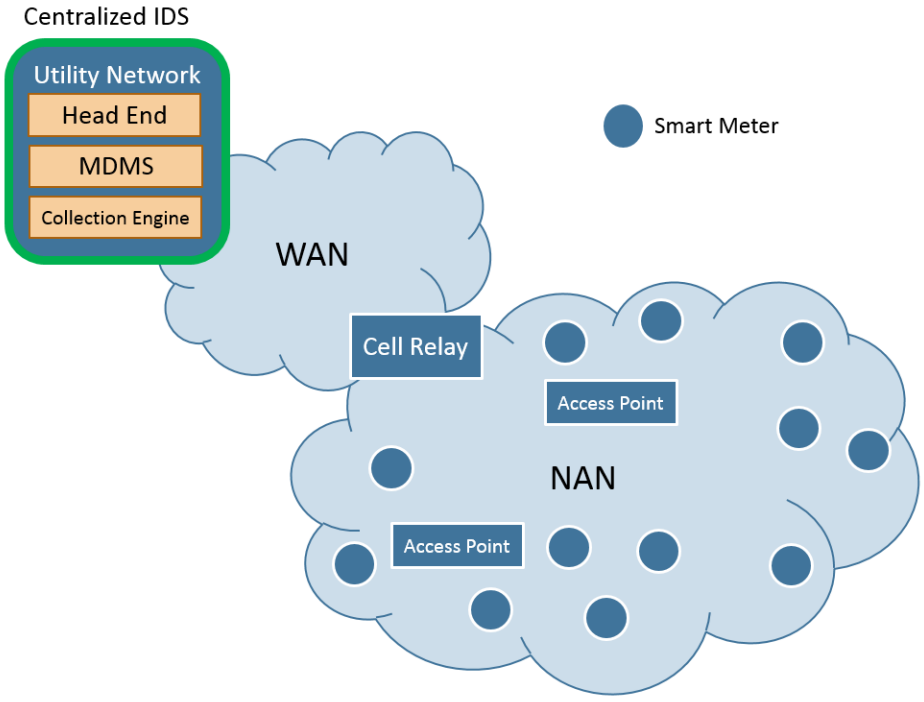

Figure 3.1: Centralized IDS Infrastructure

3.3 Embedded Infrastructure

An embedded sensor infrastructure (like the one depicted in Figure 3.2) in which IDS sensors are placed within the smart meters could be used to effectively monitor all of the traffic passing between meters as well as gain access to meter-specific information. Sensors could be placed in all meters, or alternatively, IDS sensors could be placed in a selected subset of strategic meters (as discussed by Shin et al. [32]). The latter option would minimize the number of sensors required while still providing many of the benefits of this architecture.

Since meters in the mesh network function not only as end points but also as relays, meters instrumented with sensors would have complete visibility over meter-specific information such as health reports, firmware and software integrity, and memory contents. Access to that kind of information would make it easier to detect attacks that target individual meters. A collection of meters with embedded sensors would have a holistic view of traffic flowing within the AMI network. The large number of sensors would provide redundancy to the system and increase the fault tolerance of the security monitoring system, greatly improving the accuracy and increasing the trustworthiness of generated alerts. That level of accessibility would also allow instrumented meters to be effective at analyzing the contents of messages flowing within the AMI and at detecting malicious packets.

Furthermore, embedded sensors would be effective at detecting attacks that originate from within the HAN (e.g., through a compromised smart device or appliance) or within other attached networks and devices. Such attacks, in order to propagate to other parts of the AMI, would first need to target the meter. The meter sensors would have the ability to inspect incoming messages and detect any unusual or prohibited commands before they can be executed or passed along. With such an embedded architecture, no specialized equipment or training would be required, as the sensors would be components of the smart meters. Likewise, permits would not be needed to install and run the infrastructure (beyond whatever was needed to install the AMI network itself). Thus, there would be savings in terms of both time and cost.

Attacks that are performed directly on the DCU or components besides the smart meters, however, would be missed by the meter sensors. Detection of such attacks would require data

that would be out of the embedded sensors' reach. In contrast, the centralized IDS would have access to DCU-specific data and be more effective in detecting such attacks. Also, since some portions of the AMI traffic might be encrypted, the instrumented meters would need to have access to all of the necessary decryption keys in order to properly inspect network packets passing through from other meters. Sharing of keys would need to be done through a secure communication channel and protected well at the meters, since storage of these keys on the meters would increase the impact of a compromise of one of these units. On the other hand, the centralized IDS, as noted previously, would have access to all encryption keys, giving it easy access to encrypted traffic.

It is also worth noting that most smart meters have very limited processing power, storage, and communication capabilities. Limited processing power would limit the amount of analysis that could be done at each sensor, and it would be difficult to deploy a resource-intensive sensor, as it might be a detriment to the meter's daily operations. Installing a proper sensor might require hardware upgrades that would increase the per-meter costs to the utility. While meter vendors could sell more powerful meters that could handle such intrusion detection functions, utilities might be unwilling to pay the additional price. Most utilities need to purchase millions of smart meters, so a small increment in price for each meter (e.g., a few dollars) could result in an additional investment of millions of dollars.

3.4 Dedicated Infrastructure

An alternative deployment scheme would be a dedicated sensing infrastructure, as shown in Figure 3.3, in which a small number of dedicated sensors are deployed to monitor NANs. The dedicated sensors could be used to monitor not only security events, but also the health of the network (e.g., number and type of routing error messages or retransmissions of control packets). The key advantage would be that dedicated sensors, unlike embedded sensors, deliver high availability of processing power and storage, as the dedicated devices would be more powerful than smart meters. In addition, the number of dedicated devices needed to perform tasks would be smaller than the number of smart meters that would be needed to perform the same tasks.

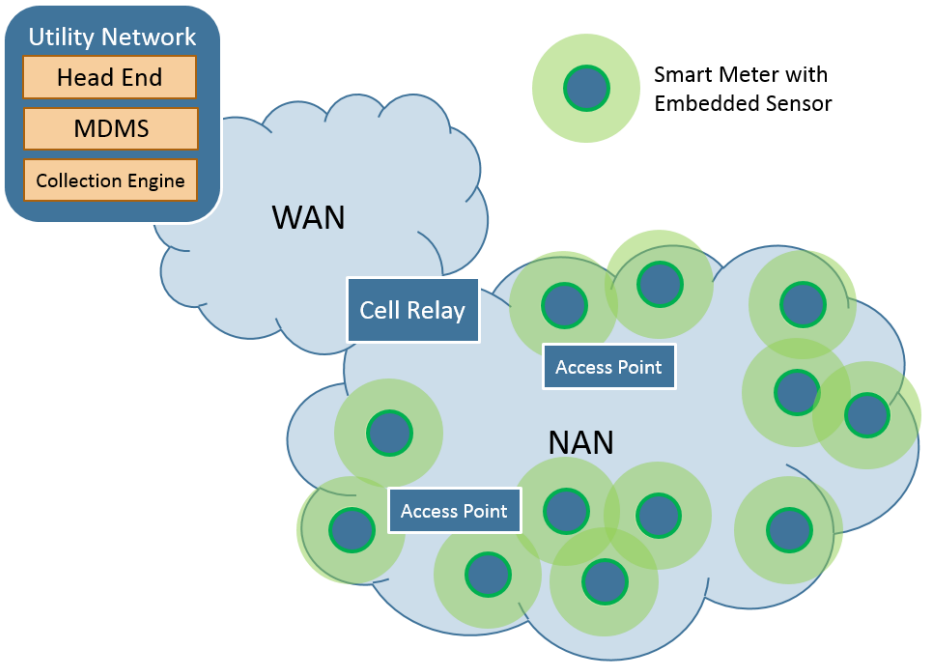

Figure 3.2: Embedded IDS Infrastructure

Dedicated infrastructure allows for the performance of much more complex IDS functions (e.g., full specification-based detection [29]) while having no impact on daily meter operations. This solution offers an interesting trade-off between network visibility and deployment cost: the number of dedicated sensors needed to cover NANs would be far smaller than the total number of meters, but the dedicated infrastructure would have smaller detection coverage.

Like the sensors in the embedded sensing infrastructure, dedicated sensors would have access to the AMI network data and be capable of monitoring all of the traffic flowing within the AMI network. Decryption keys would still need to be distributed to the sensors, but would be more manageable, since the number of dedicated sensors sharing keys would be much smaller than the number of embedded sensors that would need to share keys. Thus, the attack surface for key compromises would be greatly reduced.

However, dedicated sensors would not have access to meter-specific or DCU-specific data; therefore, they would be unable to monitor for attacks performed directly on the meters or the DCU units. Attacks originating within a HAN attempting to compromise a meter would go unnoticed until the meter started to communicate with other units. Similarly, attacks

from a smart meter targeting a device within a HAN would go undetected as well.

From a practical point of view, the cost and complexity of installation and maintenance must also be considered when deploying this type of IDS infrastructure. While smart meters have reserved sockets in which they are installed, dedicated sensors would have to be placed elsewhere (e.g., on light poles or rooftops). Such an installation would usually require a site survey, permits, renting of installation sites, and highly specialized personnel, especially if the installations were in places that were difficult to reach. Similarly, maintenance of the dedicated sensors would require that specialized personnel and equipment gain physical access to the IDS sensors. In addition, in multi-channel networks, different sets of nodes or even different pairs of nodes may communicate simultaneously on different channels, and channel selection may change on a per-packet basis. That would require that the sensor be able to decode traffic on multiple channels simultaneously. Off-the-shelf hardware with that kind of functionality is not currently available for all PHY/MAC layers in use in smart grid networks, increasing the cost of the IDS sensor.

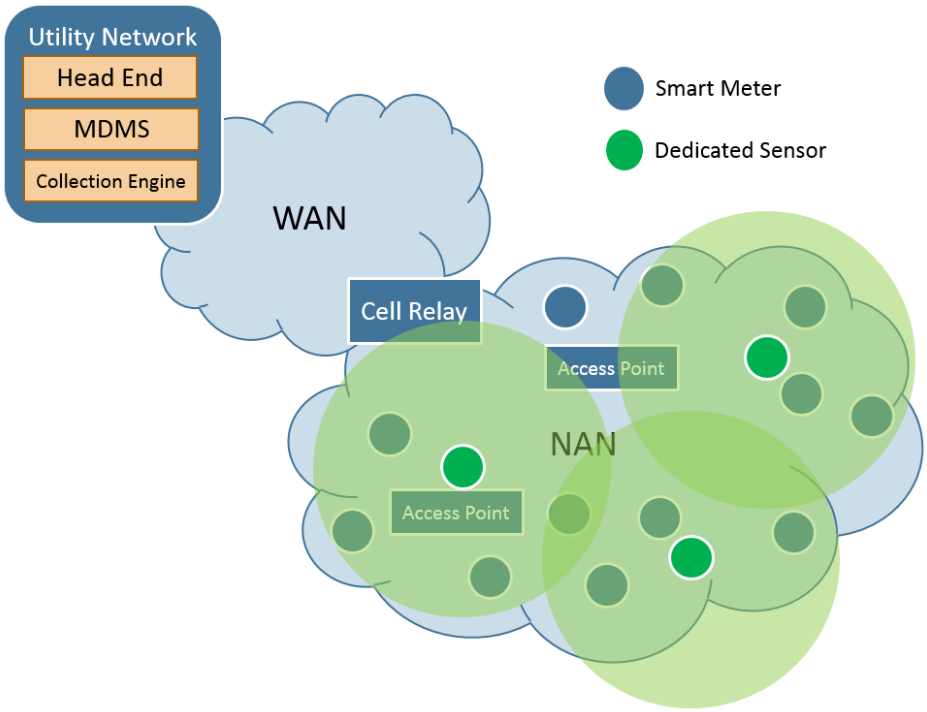

Figure 3.3: Dedicated IDS Infrastructure

3.5 Hybrid Infrastructure

A combination of the centralized sensor and embedded meter sensors might provide the widest coverage in detecting attacks. Attacks that could not be covered by the centralized IDS (e.g., attacks performed directly on the meters, or malicious packets that flow within the AMI) would be covered by the meter sensors. The meter sensors would also cover attacks that originate in the HAN. Nevertheless, it might be hard to convince meter vendors to embed the sensor capabilities, as they could push costs up where margins are already small.

Alternatively, dedicated sensors could also be used together with the centralized sensor, as shown in Figure 3.4, to monitor the traffic that flows within the AMI and manage complex IDS operations. There could be more financial incentives for security companies to build dedicated sensors and utilities to deploy them, as fewer sensors might be needed (relative to the number of meters), especially in dense urban areas. Attacks that originate from the HAN or attacks performed directly on the meters would be missed, thus reducing the monitoring coverage.

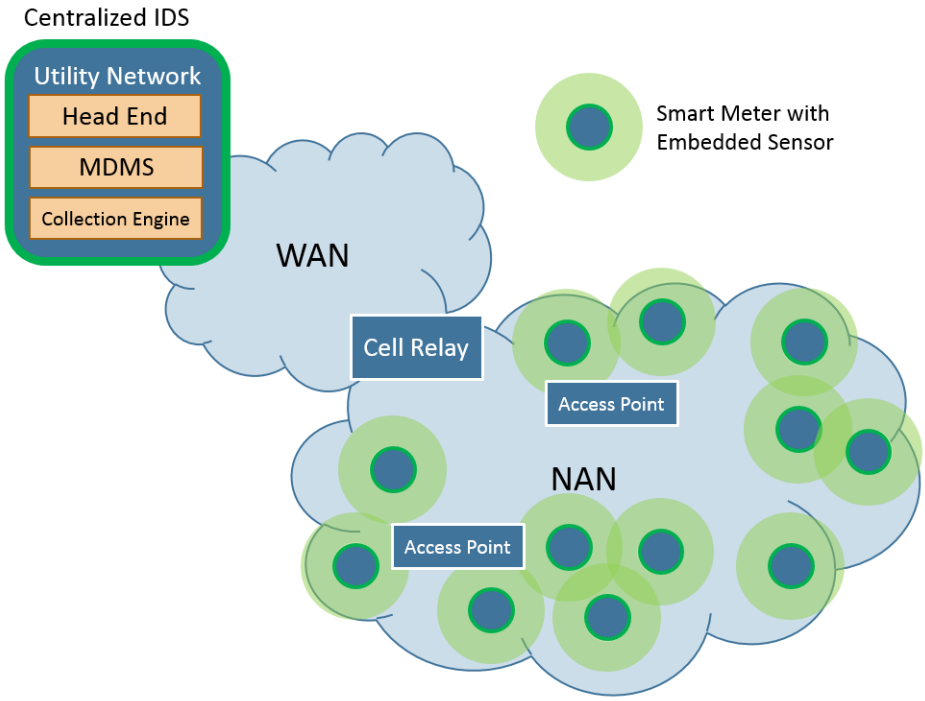

Figure 3.4: Hybrid IDS Infrastructure

3.6 Geolocation Coverage Analysis

In deciding which IDS architecture to deploy, a utility must take many factors into consideration, and there likely is not a one-size-fits-all solution. It has been suggested that a hybrid approach to intrusion detection in AMI, consisting of a centralized detection unit at the utility and distributed sensors in the field, may provide the widest detection coverage [14]. In this section, we model such a hybrid architecture from the point of view of a utility company looking to gain situational awareness over the security state of its infrastructure, while taking into consideration the need to create a business case for investing in IDSes for their AMI network. We do not perform the analysis for every possible architecture, as the goals and requirements for each deployment are unique. Rather, we provide the model as a demonstration of one process that utilities can use to select the best security monitoring solution for their AMI deployments. Because of the large disparities among AMI deployments and the complications in estimating risks, we first present our model at a high level, with the goal of providing a framework for further study.

3.6.1 Model Description

To evaluate the optimal IDS sensor density for a given AMI deployment, we performed our analysis using a stochastic model of a network of smart meters. In our model, we are able to control characteristics of the AMI deployment (e.g., sensor coverage), in addition to other aspects of the system, like IDS sensor properties (e.g., attack detection probabilities). We then performed a series of case studies in which we assigned values to those parameters to examine three realistic environments: urban, suburban, and rural AMI deployments.

The scope of our analysis is limited to modeling the reliability of a network of distributed IDS sensors by first detecting an attack at the source (e.g., the meter being attacked). The model does not model attack propagation throughout the AMI network, nor does it differentiate between different types of attacks or take into account attacks that do not target individual meters. The use of dedicated sensors offers an interesting trade-off between network visibility and deployment cost, since the number of sensors needed to sufficiently

cover AMI deployments would be far smaller than the total number of meters.

However, evaluating optimal IDS sensor density does raise questions regarding meter coverage and detection reliability, which are examined in our model. Because of state-space concerns, we assume single coverage for meters in the AMI network, though we are able to change the percentage of meters that are covered by at least one sensor. For example, in a sample deployment of 1000 meters in an urban environment, perhaps 90% of the meters are covered by at least one of 25 dedicated IDS sensors. Given the same 25 dedicated IDS sensors and a 1000-meter deployment in a rural area, perhaps only 35% of the meters are covered by the IDS sensors. What is the reliability of the first system compared to the second?

In order to model the detection reliability of a dedicated IDS sensor architecture, we first developed a model representing an AMI deployment in which the IDS sensors would be deployed. The entire model was developed in Möbius [33], a software tool for modeling the behavior of complex systems. In order to lay the groundwork for future study, our initial model is composed of two meters, which represent two types of meters in an AMI deployment. One meter lacks any coverage from an IDS sensor, while the other one is covered by at least one IDS sensor. We use these two meters to represent a large sample of meters in an AMI network.

If an attack occurs in the AMI network, we assume that all meters are equally likely to be targeted. Our models of the two meters have three primary components: sensor coverage, which describes the number of dedicated sensors covering the meter (which is 1, for the purposes of this thesis); attack targeting, which determines whether the meter is the target of an attack, given that an attack has occurred; and attack detection, which determines whether an attack is detected, given that there is coverage by at least one sensor and an attack is targeting the meter.

When an attack occurs, the attack-targeting component of the model selects one of the two types of meters to direct the attack, based on the IDS coverage discussed in Section 3.6.2. We assume that all of the meter types have the same chance of being targeted. If the attack targets a meter that lacks IDS coverage, the attack automatically goes undetected. If the attack targets a meter that is covered by a dedicated sensor, the attack detection

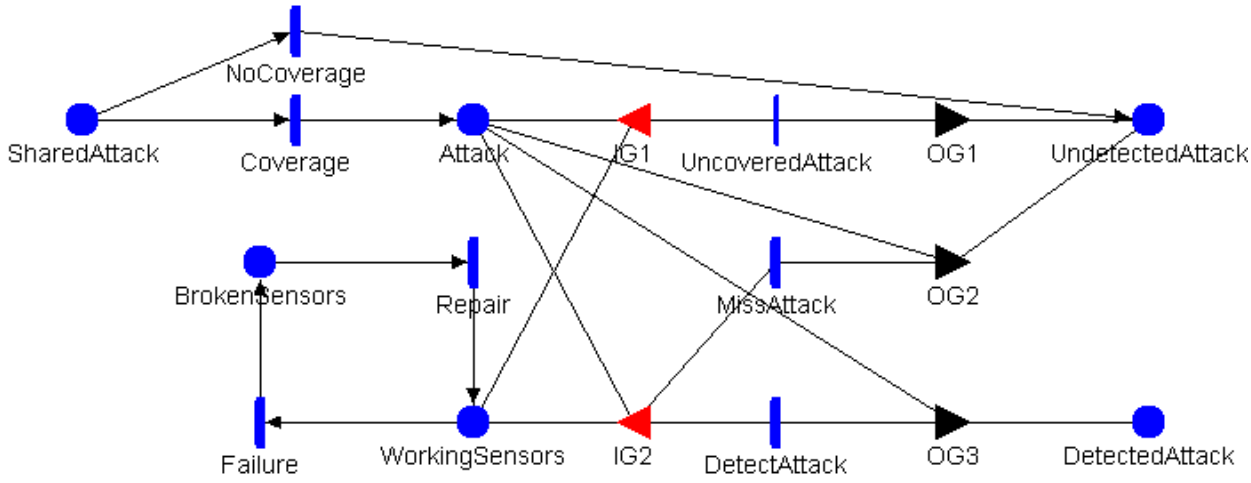

Figure 3.5: SAN Model of Basic AMI Deployment

component of the targeted meter takes over. As long as there is coverage (sensors fail with rate f_s, and are repaired with rate r_s), the attack will be detected with some detection rate r_d.

Our model, in its entirety, is shown in Figure 3.5. Additional documentation on this model can be found in Appendix A.

3.6.2 Coverage Percentage Analytical Study

We first explore IDS coverage for each of the three deployments (urban, suburban, and rural) in an analytical study. Among the various inputs required by our model, the percentage of meters that are covered by a set number of dedicated sensors p_c requires additional study. The number of meters n_m is well-known by utilities, since it matches their customer base. We can express the number of smart meters in terms of area through the use of several basic meter densities derived from U.S. Census data on household densities [34], assuming one meter per household. In particular, the densities for urban, suburban, and rural environments are 972.5, 128.7, and 7.5, respectively. Holding the total number of meters being monitored, n_m, constant at 5,000, we are looking at a coverage area of

$$Area = 5,000/Household\_Density \tag{3.1}$$

Substituting in the proper household density, we achieve $A_u = 5.14$, $A_s = 38.85$, and $A_r = 666.67$ for urban, suburban, and rural environments, respectively. Those results are reasonable, since the total area of an urban environment is smaller than the surrounding suburban area, which, in turn, is smaller than the surrounding rural area.

Coverage of a wireless network by dedicated sensors requires that the sensors be deployed such that they can overhear any packet sent by a smart meter within the AMI network area. A typical density used as a guideline for the deployment of Wi-Fi access points is in the range of 50-60 nodes per square mile, given a range of 300 feet per node [35]. We adopt that guideline for our dedicated sensor evaluation. Utilizing the same household density figures, we can express p_c in terms of the area that needs to be covered and the number of dedicated sensors n_s, as follows:

$$p_c^u = ((n_s/60) * 972.5)/n_m = 0.003242 * n_s \tag{3.2}$$

$$p_c^s = ((n_s/60) * 128.7)/n_m = 0.000429 * n_s \tag{3.3}$$

Rural environments are a special case. Assuming that meters are perfectly dispersed over a given area, we need only one sensor per meter to provide coverage. Any redundant sensors per area would be a waste and could otherwise be used elsewhere. For the rural environment case, we arrive at the following probability for rural environments:

$$p_c^r = (n_s/7.5)/n_m = 0.000027 * n_s \tag{3.4}$$

3.6.3 Experiment

In order to execute our model, we must select values for important parameters (i.e., f_s, r_s, and r_d). The rate of detection will vary widely among IDS implementations. For the purpose of our study, we assumed $r_d = 0.95$. Similarly, the failure rate of an IDS sensor will vary among hardware implementations, and the repair rate of broken sensors will vary according to utility response times. In our model, we assume that the IDS sensors have a 95% uptime, resulting in $f_s = 0.05$ and $r_s = 0.95$.

Table 3.1: Table of IDS coverages, associated mean detection percentages, and 95% confidence intervals for urban, suburban, and rural environments

# Sensors	p_c^u	μ_d^u	p_c^s	μ_d^s	p_c^r	μ_d^r
100	0.324	30.658 ± 21.259	0.043	4.057 ± 3.892	0.003	0.255 ± 0.255
200	0.648	61.315 ± 23.720	0.0868	8.114 ± 7.455	0.005	0.511 ± 0.508
300	0.973	91.972 ± 7.383	0.129	12.170 ± 10.689	0.008	0.766 ± 0.760
309	1	94.564 ± 5.141	0.133	12.535 ± 10.964	0.008	0.789 ± 0.783
400	1	94.564 ± 5.141	0.172	16.227 ± 13.594	0.011	1.021 ± 1.011
500	1	94.564 ± 5.141	0.215	20.284 ± 16.170	0.014	1.277 ± 1.260
600	1	94.564 ± 5.141	0.257	24.341 ± 18.416	0.016	1.532 ± 1.508
700	1	94.564 ± 5.141	0.300	28.397 ± 20.333	0.019	1.787 ± 1.755
800	1	94.564 ± 5.141	0.343	32.454 ± 21.921	0.022	2.043 ± 2.001
900	1	94.564 ± 5.141	0.386	36.511 ± 23.180	0.024	2.298 ± 2.245
1000	1	94.564 ± 5.141	0.429	40.568 ± 24.110	0.027	2.553 ± 2.488
1100	1	94.564 ± 5.141	0.472	44.624 ± 24.711	0.030	2.809 ± 2.730
1200	1	94.564 ± 5.141	0.515	48.681 ± 24.983	0.032	3.064 ± 2.970
1300	1	94.564 ± 5.141	0.558	52.738 ± 24.925	0.035	3.319 ± 3.209
1400	1	94.564 ± 5.141	0.601	56.795 ± 24.538	0.038	3.574 ± 3.447
1500	1	94.564 ± 5.141	0.644	60.852 ± 23.822	0.041	3.830 ± 3.683
1600	1	94.564 ± 5.141	0.686	64.908 ± 22.777	0.043	4.085 ± 3.918
1700	1	94.564 ± 5.141	0.729	68.965 ± 21.403	0.046	4.340 ± 4.152
1800	1	94.564 ± 5.141	0.772	73.022 ± 19.700	0.049	4.596 ± 4.385
1900	1	94.564 ± 5.141	0.815	77.079 ± 17.667	0.051	4.851 ± 4.616
2000	1	94.564 ± 5.141	0.858	81.135 ± 15.306	0.054	5.106 ± 4.846
2100	1	94.564 ± 5.141	0.901	85.192 ± 12.615	0.057	5.362 ± 5.074
2200	1	94.564 ± 5.141	0.944	89.160 ± 9.665	0.059	5.617 ± 5.302
2300	1	94.564 ± 5.141	0.987	93.306 ± 6.246	0.062	5.872 ± 5.528
2332	1	94.564 ± 5.141	1	94.563 ± 5.141	0.063	5.954 ± 5.600
2400	1	94.564 ± 5.141	1	94.563 ± 5.141	0.065	6.128 ± 5.752

Table 3.1 contains the results of execution of our model for a variable number of dedicated sensors, utilizing the coverage equations discussed previously and the parameters specified in the previous paragraph. We illustrate mean detection probabilities, μ_d^u, μ_d^s, and μ_d^r, for the three different environments in Figure 3.6.

3.6.4 Discussion

Our results are consistent with intuition. Increasing the number of dedicated sensors in an urban environment quickly increases the detection reliability of the distributed IDS system,

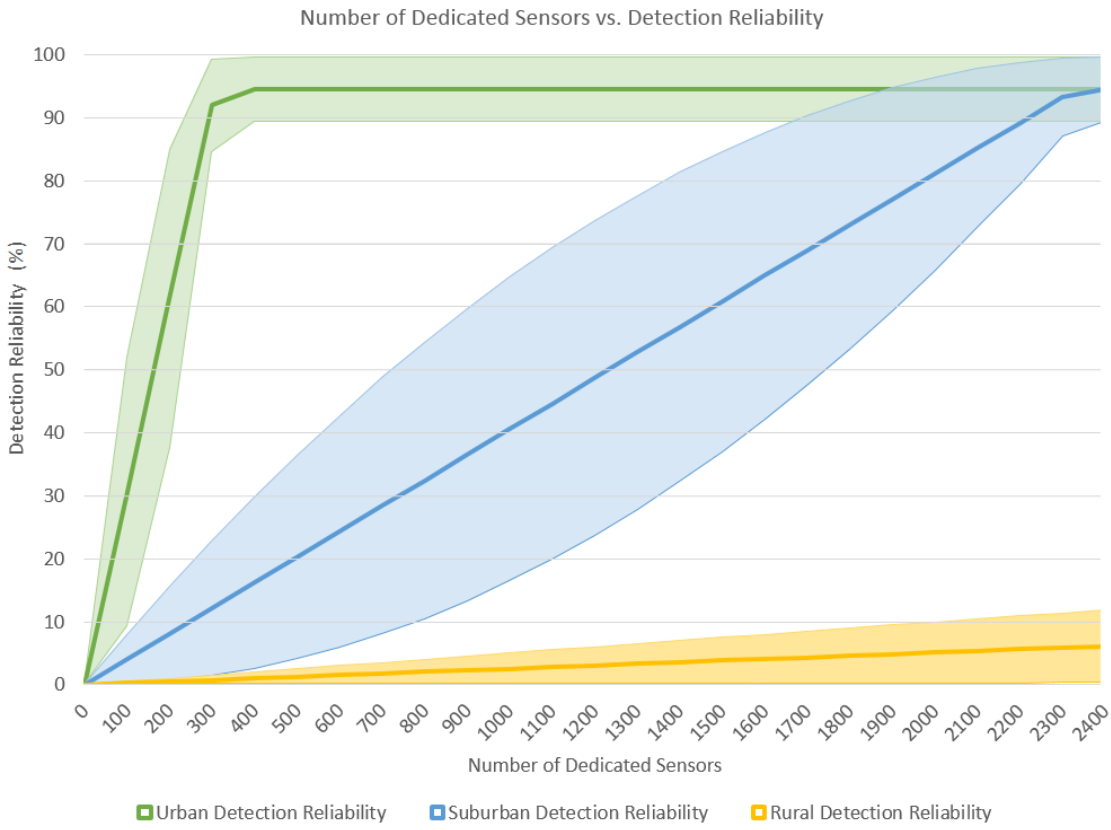

Figure 3.6: Number of Dedicated Sensors vs. Detection Reliability for Urban, Suburban, and Rural AMI Deployments

because of the high density of meters in such an environment. In a suburban deployment, a larger number of sensors is required to achieve the same detection reliability. For rural deployments, a significantly larger number of sensors is required to achieve the same detection reliability. That suggests that for rural deployments, it may be infeasible for utilities to place sensors that provide high coverage.

In fact, the business case for deploying IDSes in a rural environment maybe even worse than what is suggested by our results. While we assume that there is one sensor per meter in rural AMI deployments, there may in fact be many more. Some rural AMI networks consist of a large chain of smart meters forming a line. Deployment of dedicated IDS devices to monitor such a network would require at least one dedicated IDS sensor for every two smart meters, if not one for every smart meter. In contrast, in urban deployments, we have seen sensor nodes with several hundred (or even a thousand) smart meters within the

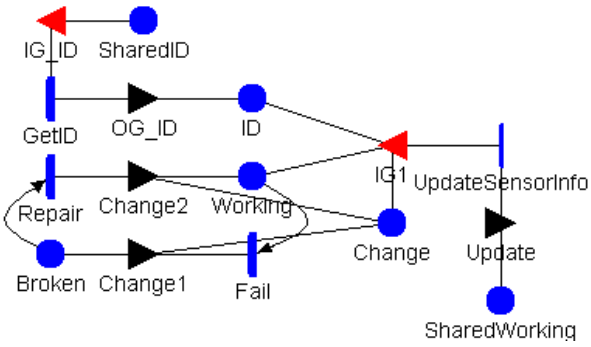

Figure 3.7: SAN Model of an IDS Sensor System

same wireless communication range; therefore, a single dedicated IDS device can potentially monitor hundreds of smart meters, making a dedicated infrastructure cost-effective in large, high-density deployments. Our results seem to suggest similar cost-effectiveness for suburban environments as well.

Additionally, the detection reliability does not surpass the sensor detection rate of 0.95 regardless of the AMI deployment environment. That shows empirically that the detection rate of an IDS system, regardless of sensor reliability or the size of the deployment, is limited to the quality of the IDS algorithm being utilized.

3.7 Detection Coverage Analysis

While our first model offered a high-level insight into how sensor coverage varies by deployment environment, it did not fully represent how sensors provide coverage to meters in a dedicated IDS infrastructure. In a real-world deployment, per-meter coverage is not always attained. While some meters may be monitored only by one sensor, depending on coverage, there may be some meters that are covered by more than one sensor and others that lack any coverage at all. For example, in a sample deployment of 1000 meters and 25 sensors, 90% of the meters may be covered by at least one sensor, 40% by at least 2 sensors, and 15% by more than 2 sensors. How would the addition of another optimally placed sensor affect reliability of the system, and hence the probability of detecting or preventing a specific failure scenario?

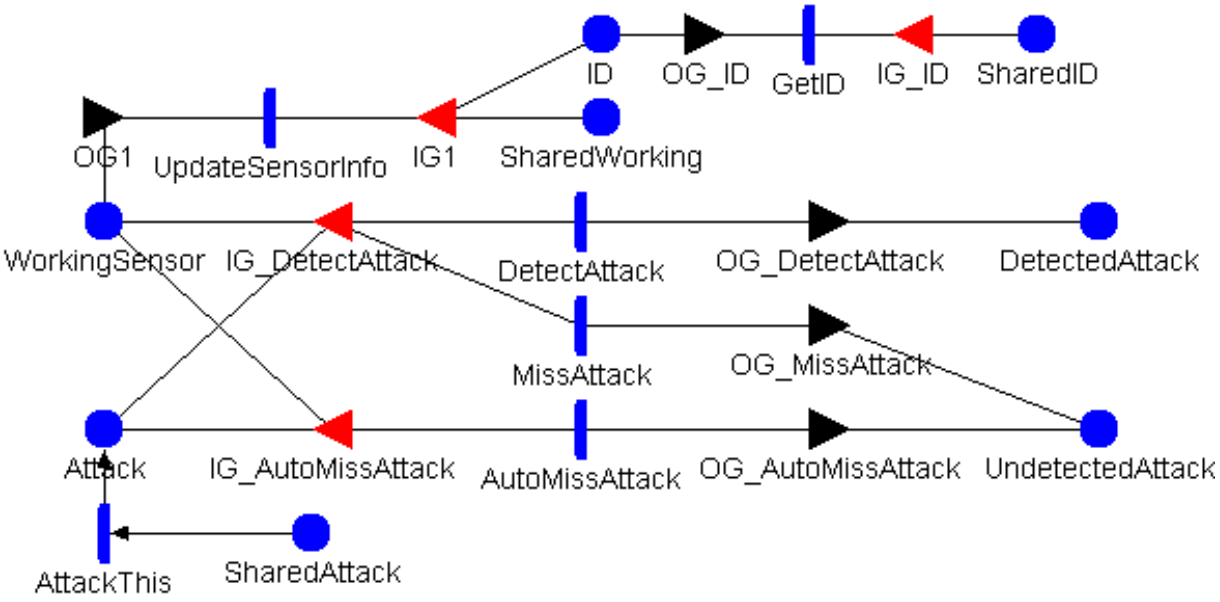

Figure 3.8: SAN Model of an AMI Meter

3.7.1 Model Description

To evaluate the optimal IDS sensor density for an environment with varying degrees of coverage, we again performed our analysis using stochastic models, but this time with a more advanced model of a network of smart meters that can have varying coverage. We also applied the model to three of the five case studies discussed in Section 2.3: DDoS, remote disconnect, and stealing of customer information. We modeled the failure scenarios using Adversary View Security Evaluation (ADVISE) models, in which we could control attributes of the adversaries such as specific technical skill levels, the appeal of reaching specific goals, and willingness to perform certain actions despite the associated detection probabilities. Our improved model shares some of the limitations of the first. In particular, it is still limited to modeling the reliability of a network of distributed sensors attempting to detect an attack at the source (e.g., the meter being attacked).

In this model, we represented each sensor with a single token in a simple cycle rotating between two states, working and broken, at certain exponential rates, repair rate r_r and failure rate r_f (Figure 3.7). We represented each meter as a more complex model in which a token representing an occurring attack would result in either the detected state or the

undetected state (Figure 3.8). That final state would be determined by whether or not the specific meter was covered by at least one operational sensor at the time of the attack and the probability that the sensor could detect the given attack.

In order to model each smart meter system failure scenario, we next developed a model with the ADVISE formalism for each given scenario. In the case of a DDoS attack, we determined that meters must first be compromised before executing the attack (Figure 3.9). In order to compromise a meter, the adversary first had to install malware on a meter by taking advantage of physical access to the meter, insider privileges, or known network vulnerabilities. The detection probability and execution time of the malware installation were dependent on the access to the meter, the adversary's level of DDoS skill, and the number of sensors covering the targeted meter. The success rate of the DDoS attack depended on the number of compromised meters.

Second, to handle remote disconnects, we determined that in order to send a remote disconnect command, the adversary had to have access to a compromised meter and had to have specific information about the target (e.g., IP address) (Figure 3.10). In order to acquire the specific information about the target, the adversary had to have access to a compromised meter. The adversary could acquire that access by installing malware on a given meter (just as in the DDoS scenario). The detection probability and execution time of the malware installation would depend on the access to the meter and the number of sensors covering the targeted meter. The detection probability, execution time, and success probability of the specific meter information acquisition and remote disconnect command transmission depend on the adversary's skill level pertaining to the wireless meter protocol.

Third, in order to steal a customer's information, the adversary had to know the appropriate decryption keys and have access to the encrypted messages sent from said customer (Figure 3.11). Obtaining the decryption keys was possible through brute-forcing of the cryptographic system of the AMI. The encrypted messages sent from the customer could be obtained either by compromising a meter with malware or by using an RF sniffer on the same frequencies of the AMI air traffic. The detection probability and execution time needed to brute-force the crypto system depended on the access to the meter, the number of sensors covering the targeted meter, and the adversary's cryptography skills. The detection proba-

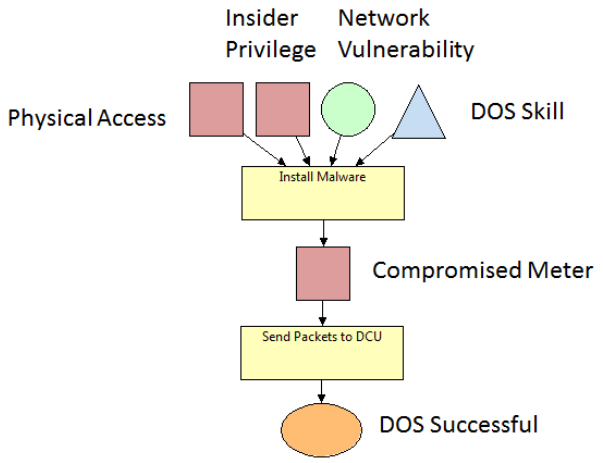

Figure 3.9: ADVISE Model of DDoS Attack Scenario

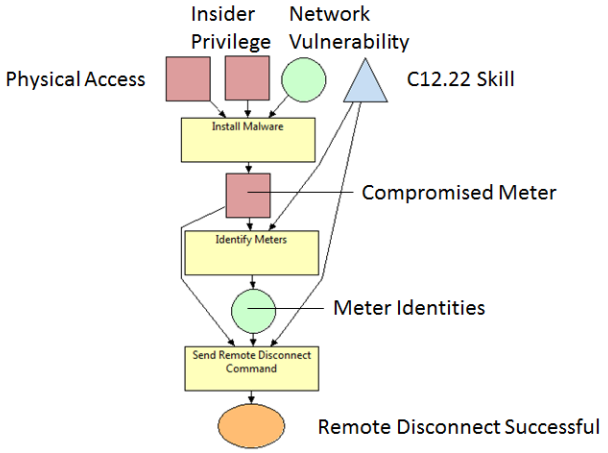

Figure 3.10: ADVISE Model of Remote Disconnect Attack Scenario

bility and execution time of installing malware on a meter depend on the access medium to the meter, the number of sensors covering the targeted meter, and the network-sniffing skills of the adversary. The execution time needed to eavesdrop on the AMI traffic depended only on the adversary's network-sniffing skills.

Additional documentation on this model can be found in Appendix A.

3.7.2 Adversary Classification

Because of the complexity of the multitude of possible adversary backgrounds, we parameterized the adversaries as the most likely candidates for each given scenario in terms of skill

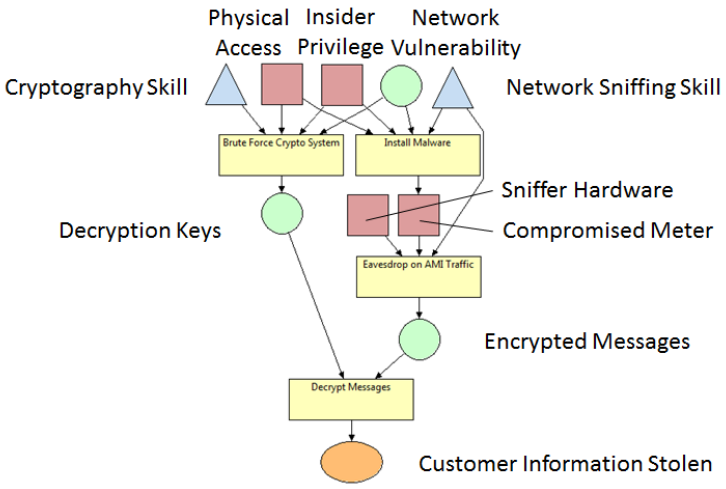

Figure 3.11: ADVISE Model of Information Theft Attack Scenario

levels, initial accesses, initial knowledge, willingness to take risks, and motivation to reach a final goal. For example, we classified the DDoS adversary as an activist group, or a group that is willing to disrupt order. The group has low to medium technical skills, low access levels and initial knowledge, high willingness to take risks, and medium to high motivation to reach a final goal. Next, we classified the remote disconnect adversary as a corporation, an adversary that is willing to maliciously disconnect another corporation from power in order to ruin a product or just cause general disruption. That adversary has high skill levels, low to medium initial access, medium to high initial knowledge, low willingness to take risks, and medium motivation to reach the final goal. Finally, we classified the adversary most likely to steal customer information as a lone burglar, an adversary who may use the information to determine whether an occupant is away in order to plan a robbery. That adversary has low skill levels, low initial access and knowledge, high willingness to take risks, and high motivation to reach the final goal.

3.7.3 Experiment

To analyze the effectiveness of additional dedicated IDS sensors, we set up four model architectures: one in which an attack always occurs and three in which the three previously discussed attack scenarios occur. For each of the models, we used a fixed value of 20 meters

Table 3.2: Table of IDS coverages, attacker willingnesses, and 95% confidence intervals for general, denial of service, remote disconnect, and privacy attacks

Sensor Density	p_d^{Attack}	p_d^{DOS}	W_{DOS}	p_d^{Rem}	W_{Steal}	p_d^{Steal}	W_{Steal}
0.000	0	0.040 ± 0.039	1.000	0	1.000	0.020 ± 0.028	1.000
0.250	0.200 ± 0.079	0.280 ± 0.088	1.000	0.300 ± 0.090	1.000	0.320 ± 0.092	1.000
0.500	0.470 ± 0.098	0.480 ± 0.098	1.000	0.400 ± 0.097	1.000	0.470 ± 0.098	1.000
0.750	0.730 ± 0.087	0.800 ± 0.079	1.000	0	0.000	0.720 ± 0.088	1.000
1.000	0.900 ± 0.059	0.940 ± 0.064	1.000	0	0.000	0.940 ± 0.047	1.000
1.250	0.920 ± 0.053	0.870 ± 0.066	1.000	0	0.000	0.940 ± 0.047	1.000
1.500	0.950 ± 0.043	0.930 ± 0.050	1.000	0	0.000	0.900 ± 0.059	1.000
1.750	0.970 ± 0.034	0.900 ± 0.059	1.000	0	0.000	0.960 ± 0.039	1.000
2.000	0.920 ± 0.053	0.930 ± 0.050	1.000	0	0.000	0.970 ± 0.034	1.000
2.250	0.940 ± 0.047	0.940 ± 0.047	1.000	0	0.000	0.940 ± 0.047	1.000
2.500	0.960 ± 0.039	0.930 ± 0.050	1.000	0	0.000	0.980 ± 0.028	1.000
2.750	0.930 ± 0.050	0.960 ± 0.039	1.000	0	0.000	1.000 ± 0.000	1.000
3.000	0.940 ± 0.047	0.940 ± 0.047	1.000	0	0.000	1.000 ± 0.000	1.000
3.250	0.990 ± 0.020	0.910 ± 0.056	1.000	0	0.000	1.000 ± 0.000	1.000
3.500	0.940 ± 0.047	0.970 ± 0.034	1.000	0	0.000	1.000 ± 0.000	1.000
3.750	0.960 ± 0.039	0.950 ± 0.043	1.000	0	0.000	1.000 ± 0.000	1.000

as we varied the number of dedicated IDS sensors, hence also varying the sensor coverage density. We assumed that optimal placement of the sensors allowed one sensor to cover five meters, and that multiple coverage occurred only if all meters had similar coverage. For example, four sensors could cover all 20 meters once, and the addition of a fifth sensor would result in double coverage for five of the meters. We ran each experiment 100 times.

For each of the model architectures, we analyze three trends: the probability that the attack is detected, the probability that the attack is not detected, and the probability that the attacker is unwilling to perform the attack. The three trends allow us to use certain thresholds related to the value of preventing a successful attack to determine the optimal number of sensors necessary to either detect or deter an attack.

Table 3.2 contains the results of execution of our model for variable sensor coverage density, utilizing the coverage described previously. We illustrate the detection probabilities and confidence intervals for the three different failure scenarios—denial of service, remote disconnect, and information stealing—in Figure 3.12, Figure 3.13, and Figure 3.14, respectively. In each of the figures, the red shading indicates the attacker's willingness to carry

out the attack. That can also be thought of as a measure of deterrence. The shaded regions indicate that the attacker is willing to carry out the attack. Unshaded regions indicate that the attacker finds the risk to be too high. Figure 3.15 shows the detection reliability for a general attack, and is used for comparison.

3.7.4 Results

As expected, when an attack always occurred, initial additional sensors drastically increased the probability that an attack would be detected. However, once the sensor density reached 1.5, the probability of detection leveled out around 0.95. Additional sensors after that point seemed to be in excess.

The DDoS attack showed slightly different results. As expected, the probability of the detection of the attack increased with additional sensors. However, in this model, the adversary was not likely to give up the attack. If the detection probability of an attack was greater than the risk the adversary was willing to take (0.8), then the adversary would decrease the probability of detection by acquiring more compromised meters. Since the adversary would only take these steps only if necessary, he or she could underestimate the probability of detection. Although that would increase the probability of detection for a certain number of meters, this effect appears to be negligible in our data.

The trends of the remote disconnect attack were heavily governed by the traits of the adversary: specifically, the lack of willingness to take risks (0.5). Since the adversary was a corporation that could face serious consequences for getting caught performing an attack, it was willing to commit to the attack only for the first several sensors. After that, the attack was completely deterred.

The results for the attack stealing customer information most closely match the results for an attack always occurring. The similarity is due to the adversary's high willingness to take risks. Although the adversary had lower technical skills and knowledge, he or she still performed the attack almost every time.

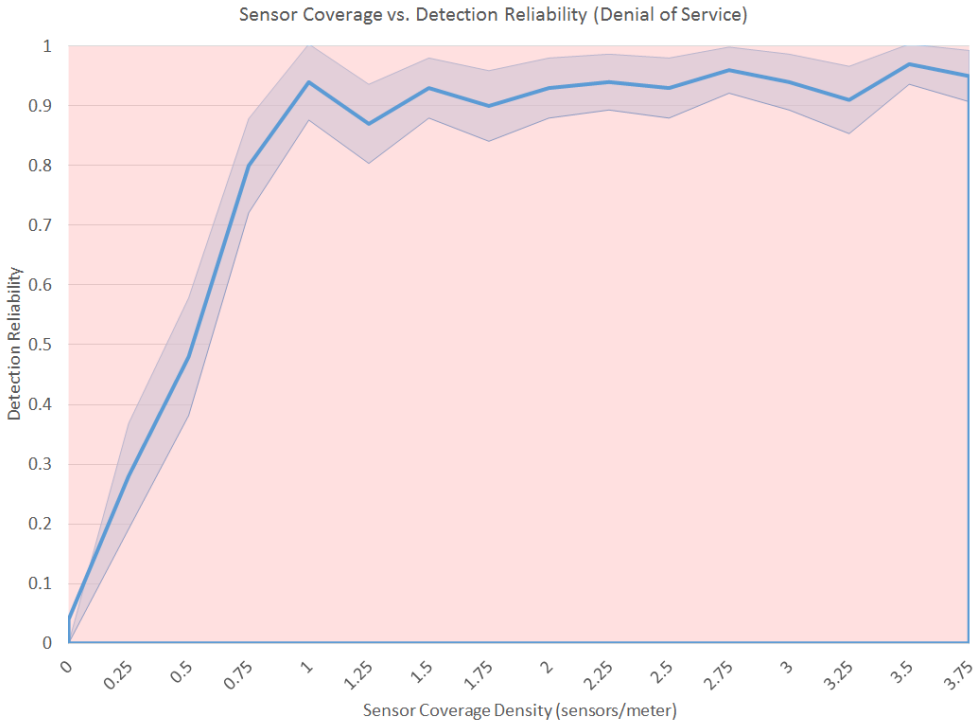

Figure 3.12: Sensor Coverage vs. Detection Reliability (DDoS)

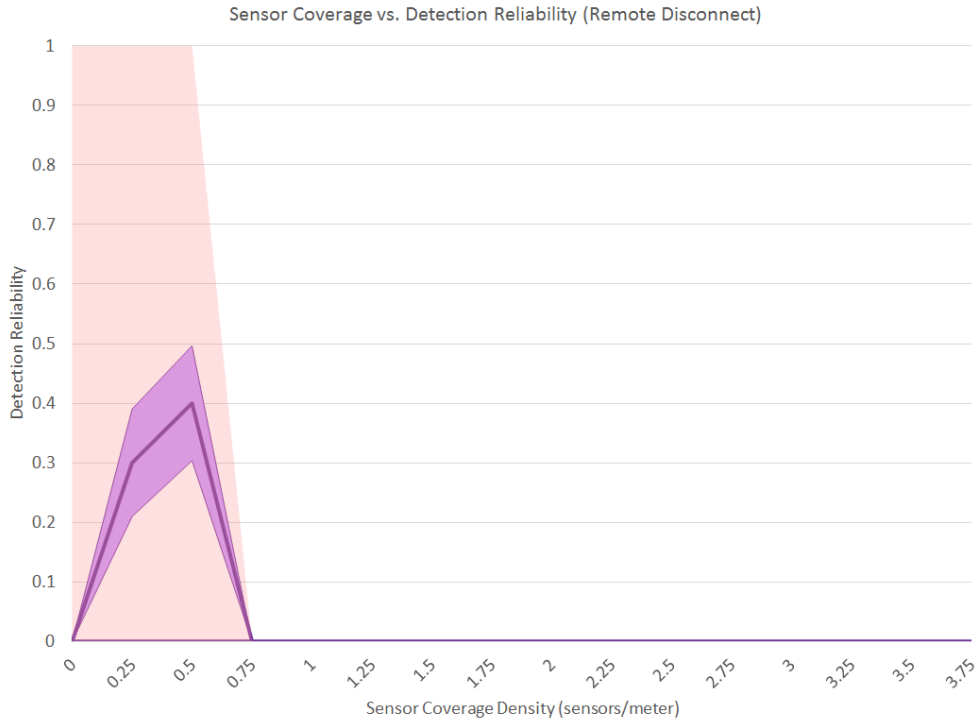

Figure 3.13: Sensor Coverage vs. Detection Reliability (Remote Disconnect)

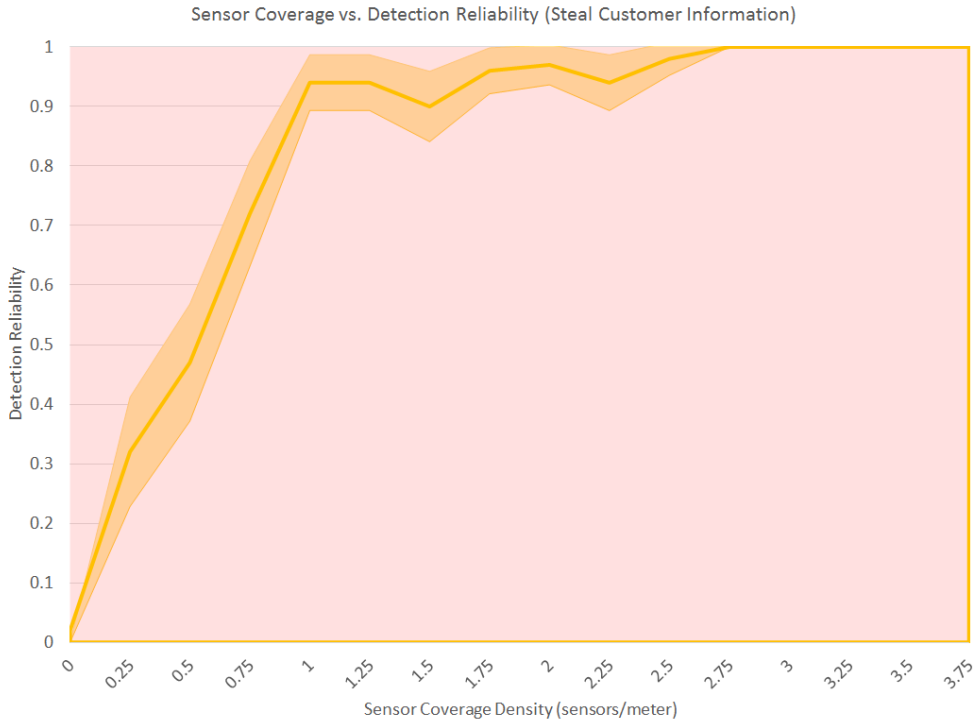

Figure 3.14: Sensor Coverage vs. Detection Reliability (Information)

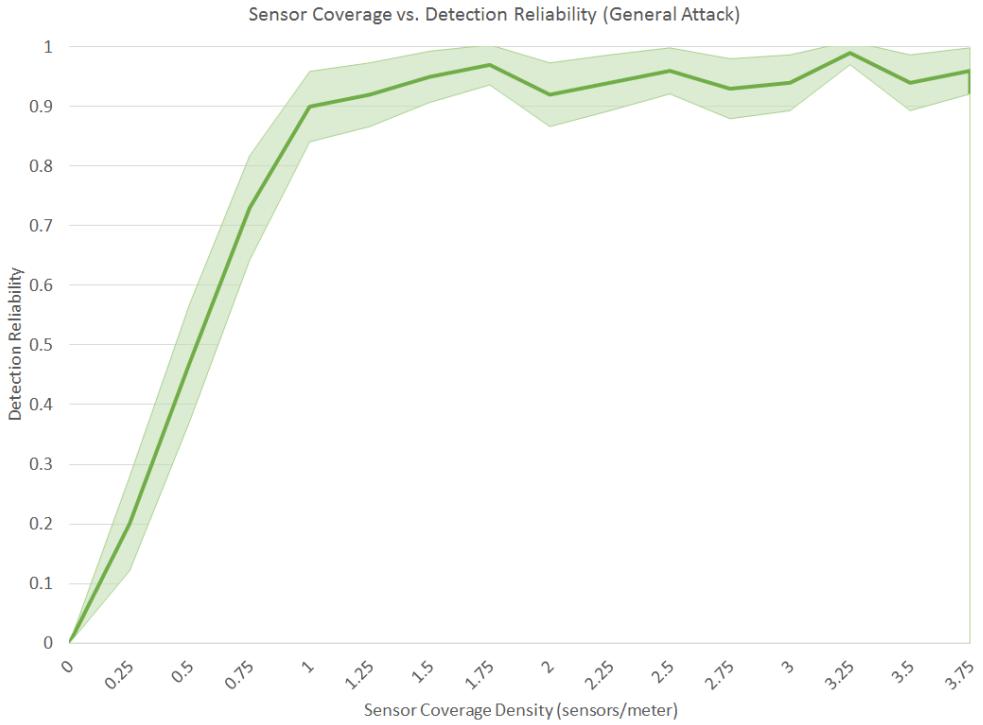

Figure 3.15: Sensor Coverage vs. Detection Reliability (General Attack)

3.7.5 Discussion

In nearly every case, increasing the density of dedicated sensors increased the detection reliability of the distributed IDS system. The reasons were the increased reliability of the sensors and the greater coverage that they provided. We can also gather from our results that while additional sensors may increase detection reliability, in all scenarios (except remote disconnect), we hit a peak a sensor density at around 1. That was the first threshold that we were interested in discovering. Further study is needed to determine whether that density can be considered optimal, but it does suggest that the cost of additional sensors to increase coverage may not be justified. As sensor coverage density approaches 1, we see a roughly linear increase in detection probability. That suggests that it may be easier than originally thought for utilities to determine the number of sensors needed to achieve a desired detection rate.

Figure 3.13 shows an interesting result. At a sensor density of 0.75, the attacker willingness drops to zero, and with it the detection probability. That illustrates a key point in attack deterrence and reflects the second threshold that we were interested in finding. Some security measures, like dedicated IDS sensors, are visible to the attacker. An attacker could, for example, drive around a neighborhood and count the number of IDS sensors. If the attacker feels there is sufficient security and fears being detected, the attack may never occur. It is important to note that while the attack detection probability drops to zero, this is expected, as there are no attacks to be detected. Further analysis needs to be done to vary additional attack parameters to detect interesting scenarios like the one displayed in Figure 3.13.

CHAPTER 4

KEY MANAGEMENT CONSIDERATIONS

While we mentioned key management concerns in Chapter 3, our analysis thus far has largely assumed that the IDS architectures have full access to the traffic being monitored. However, that may not always be the case. Many of the components in AMI networks, like the smart meters and DCUs, are capable of communicating with each other with data encryption. Depending on how monitoring solutions are integrated with the AMI, that can cause complications for both the security of the system and the reliability of the detection. In this chapter, we will explore requirements for key management in the context of the operations that would be required in a proper deployment, and then propose ways of ensuring efficient and reliable detection on AMI deployments that use encrypted traffic taking into account those requirements. Finally, we discuss possible next steps and future work in designing an IDS-friendly key management system for AMI.

4.1 Main Key Operations

Secure communication between among devices in AMI is essential in the overall security of the network. While some of these communications might not contain compromising information, sensitive customer information and proprietary utility data are frequently being sent over the system. To address these concerns, the ANSI C12 standard [36] was developed to specify requirements for secure communication among devices in AMI networks. According to the C12.22 standard, messages can be transported in three forms: cleartext, authenticated cleartext, and authenticated ciphertext. However, it is up to device manufacturers and utilities to set these parameters in their deployments. In order to support the third option of transporting messages in authenticated ciphertext, certain mechanisms must be in

place. This includes generating secure encryption keys, ensuring proper communication of encryption keys to smart devices, supporting protected storage of encryption keys on smart devices, and mechanisms to support operations like revoking compromised keys and rekeying large numbers of devices on demand.

These and other factors must be considered when choosing between key management schemes. In the remainder of this chapter, we explore different ways of integrating IDSes with key management systems that meet the requirements for AMI.

4.2 Sharing Keys with Sensors

One solution would be to share the decryption keys with the IDS sensors. The IDS sensors can be given all of the keys they need to decrypt traffic on the network. With all traffic still encrypted between the meters and the utility, the sensors would be able to decrypt traffic as needed and use any IDS algorithms for the purpose of detection. That might lower the computation and power requirements of the sensor, making the embedded infrastructure possible and allowing more advanced operations on dedicated sensors. Communication of decryption keys to meters or sensors must be done over a secure channel, and they must be properly stored. There must also be mechanisms in place to revoke and update the set of keys as needed.

In the rest of this section, we will look in more detail at considerations in regard to key management and sharing keys for each of the IDS deployment architectures discussed in Chapter 3.

Centralized Infrastructure

The main advantage of a centralized IDS deployment is that the sensor will be able to perform a full analysis on encrypted traffic flowing into and out of the utility. Because of its proximity to the certificate authority, a centralized sensor would have complete access to all of the necessary cryptographic keys needed for decrypting traffic. Also, being housed within the utility data center ensures that the keys are still protected, assuming the data center is not breached.

Embedded Infrastructure

For an embedded IDS deployment, it would be more difficult to manage decryption keys. In order to enable effective detection, the sensors would need to have access to all of the decryption keys needed to monitor traffic flowing through the meter, including traffic that is being relayed to other nodes in the network. Determining which keys need to be shared might be cumbersome, especially if the topology of the network changes. Sharing of keys would need to be done through a secure communication channel, and they keys would need to be stored on the meters. We showed before that meters do not have much physical protection, so the keys would have to be properly stored, as improper storage would result in a higher impact if a unit were compromised by an adversary.

Dedicated Infrastructure

Dedicated IDS deployments would share some of the same key management problems. While there would be fewer sensors to distribute keys to, the necessary decryption keys would still need to be securely communicated and stored on the dedicated sensors. Distribution would be more manageable, meaning that the added vulnerability of storing keys in the field would be reduced, but still not eliminated.

Hybrid Infrastructure

A hybrid approach would share the same benefits and risks of the individual architectures. For instance, a deployment consisting of a centralized sensor with embedded sensors in the field would have the same distribution and storage concerns as the embedded infrastructure.

4.3 Traffic Classification

Another option would be to monitor encrypted traffic in the AMI. In that case, all traffic between meters and the utility would be encrypted, and the IDS sensors would not have access to any decryption keys. However, much work still needs to be done to apply any of the existing monitoring techniques to AMI protocols. We must also consider the power and

computational restrictions on the embedded sensors. Monitoring of encrypted traffic poses a great problem in AMIs, but many efforts have looked at this issue in other contexts. Much of the work has focused on classifying traffic.

For instance, [37] provides an early traffic analysis technique using Bayesian analysis. Using known data flows for training, the authors were able to achieve 95% classification accuracy. [38] extends the classification to a broad set of protocols with varying degrees of reliability (64%—100%, depending on protocol). Their technique requires human interaction, as rules must be created for each application (e.g., imap) based on a set of indicators. [39] takes a different approach and aims to classify traffic by observing and identifying patterns in flows generated from different hosts. Those authors were also able to achieve 95% accuracy, but without the need for any training data; they only require access to the traffic flows collected from the network. Those early techniques were promising, but were not flexible enough in live environments, nor were they as effective in cases where adversaries were deliberately trying to evade detection.

On-the-fly techniques for classifying traffic flows were first explored in [40] and extended in [41] and [42]. Bernaille et al. analyzed sample traffic flows to cluster them into classes that shared common behaviors. Then, during the online classification phase, they examined the first five packets for similar behavior. The technique achieved 98% classification accuracy. [42] extended the techniques to work on encrypted traffic flows like HTTPS. Still, techniques to circumvent detection were possible, like padding of packet payloads. [43] also looks at on-the-fly decryption of traffic that uses identity-based encryption. Those authors were able to prevent MITM attacks, but there are concerns about practicality and performance.

[44] considers a set of features to apply machine learning to a method similar to the one in [37]. The machine-learning algorithms use the most important features in the classification. Statistical methods have been used as well, with both [45] and [46] using this technique. [45] uses training data to develop rule sets used to classify traffic flows, while [46] uses statistical analysis with existing signature-based methods to classify SSL/TLS payloads with 99% accuracy.

Other researchers have focused solely on encrypted traffic. For example, [47] takes the unique approach of looking at round-trip times in communications. Those authors developed

a methodology for estimating round-trip times along specific network paths, which enables them to monitor encrypted traffic flows. While the technique does not offer insight into packet payloads, it might be able to detect MITM attacks.

[48] proposes a centralized IDS scheme that ensures that all traffic is routed through a centralized sensor as well as to the intended recipient. While the authors demonstrated that the IDS evasion rate was very low (0.98%), their method relies on all communications being sent to the sensor, which, for a large AMI deployment, would likely cause bottlenecks in communications. They expand on that work in [49], in which they developed a protocol that facilitates routing to the sensor. [50] uses a method based on deep-packet inspection and takes advantage of vulnerabilities in encrypted traffic to detect peer-to-peer traffic with a 96% detection accuracy. The method is, however, protocol-specific, and is not easily adaptable to AMI communication protocols.

Still, the above techniques are useful only in classifying network traffic. Perhaps most interesting is the work of Koch et al. [51], who present a scheme that allows for detection of commands within encrypted traffic. Their method does not involve modifying the protocol, so existing AMI communication protocols can be used, and it uses statistical analysis to identify the commands being sent. That type of monitoring would be most beneficial to sensors deployed in AMI networks.

4.4 Partial or Selective Encryption

A third option would be to selectively encrypt communications. For example, a message containing personal customer information could be signed and encrypted before being sent back to the utility, but other non-identifying pieces of information could be sent in the clear, having only been signed by the source. That would allow both the dedicated and embedded infrastructures to monitor the majority of traffic flows in the AMI network, while protecting customer information from being sent out in the clear. With that approach, the IDS sensors would not have access to decryption keys, protecting the system from increased risk if a meter or sensor were to be compromised. However, it would not be possible to monitor a message containing encrypted traffic until it reached the utility data center, where it could

be safely decrypted by a centralized sensor located there. This solution could be enhanced by including a mechanism to monitor the ratio of encrypted messages to those sent in the clear. That would help to prevent attackers from hiding their actions in encrypted traffic without raising any red flags.

Yet another possibility would be to strategically deploy decryption keys to the embedded and dedicated sensors based on the traffic that is expected to pass through them. That would limit not only the number of keys on each sensor but also the impact if a meter or sensor were to become compromised. The strategic deployment could be accomplished in either the fully encrypted or selectively encrypted environments previously discussed. Secure communication and storage of the keys would remain crucial.

4.5 Discussion

For the same reasons it is difficult to choose among IDS deployment options, it is difficult to strongly suggest the proper way to handle key management. In some situations, the answer is clear. For example, when a centralized IDS infrastructure is being used, we have already seen how key management can easily be accomplished in a safe and secure manner. However, when considering the embedded, dedicated, or hybrid infrastructures, for which we have already demonstrated supporting evidence for their benefits over a centralized approach, what is the optimal configuration?

Unfortunately, this question and others like it are outside the scope of this thesis. While there has been significant progress by other researchers looking at the approaches discussed in this chapter, further work must be accomplished in order to make these key management schemes AMI- and IDS-friendly. In particular, the best solution should be scalable, so that it can accommodate the size and magnitude of AMI deployments, as well as practical, so that it can integrate well with existing solutions. Finally, as utilities are looking to increase their situational awareness by adding security mechanisms like IDSes, selecting a key management scheme that allows IDSes to effectively monitor traffic is equally important.

CHAPTER 5

CONCLUSION

We have presented a comprehensive analysis of the possible IDS architectures for AMIs while taking into account the unique environment in which AMI networks operate.

To accomplish that, we first had to understand the threat landscape in AMIs. Our analysis of attacker motivations and possible techniques that can be used to achieve the corresponding attacker goals led to a thorough set of failure scenarios. From the failure scenarios, we were able to extract specific pieces of information that would be required to detect them effectively.

That information guided our assessment of four different IDS architectures; centralized, embedded, dedicated, and hybrid. We were able to consider the benefits and shortcomings of each and suggest that a hybrid approach, in which a centralized and dedicated sensor infrastructure is used, provides the widest coverage in monitoring of attacks. With that in mind, we modeled such an architecture in Möbius to show the process by which a utility may weigh architecture options before making an investment.

Finally, we explored various ways of deploying an IDS architecture in an AMI network that has encryption, and we proposed several possible ways of monitoring such traffic.

Those contributions will help provide utilities with the information they need to deploy IDSes in their smart grid deployments, enable effective security-monitoring solutions in AMI deployments across the country, and lay a proper security foundation for future power grids.

APPENDIX A

MODEL DOCUMENTATION

This appendix contains additional documentation for the models described in Section 3.6 and Section 3.7.

A.1 Basic Meter Model

As described in Section 3.6, the basic meter model consists of a SAN model of a smart meter. Table A.1 contains the initial markings for the places for this model (Figure 3.5). Table A.2 and Table A.3 contain the enabling predicates and completion functions for the input and output gates for the same figure, respectively. Table A.4 contains the reward functions for the performance variables used in the experiment and Table A.5 contains rates used in the execution of our model.

A.2 Expanded Meter Model

In the expanded model, there are three components: the AMI meter model, the IDS sensor system model, and the ADVISE adversary model. Table A.6 contains the initial markings for the places in the AMI meter model (Figure 3.8). Table A.7 and Table A.8 contain the enabling predicates and completion functions for the input and output gates for the same model, respectively. Similarly, Table A.9 contains the initial markings for the places in the sensor system model (Figure 3.7), and Table A.10 and Table A.11 contain the enabling predicates and completion functions for the input and output gates, respectively.

Table A.1: Initial markings for places in meter SAN model

Place Name	*Initial Marking*
Attack	0
BrokenSensors	0
DetectedAttack	0
SharedAttack	*AttackCount*
UndetectedAttack	0
WorkingSensors	1

Table A.2: Enabling predicates for input gates in meter SAN model

Input Gate	*Predicate*
IG1	$return\ (MARK(Attack)\ \&\&\ !MARK(WorkingSensors))$;
IG2	$return\ (MARK(Attack)\ \&\&\ MARK(WorkingSensors))$;

Table A.3: Completion functions for output gates in meter SAN model

Output Gate	*Function*
OG1	$MARK(Attack)--$; $MARK(UndetectedAttack)++$;
OG2	$MARK(Attack)--$; $MARK(UndetectedAttack)++$;
OG3	$MARK(Attack)--$; $MARK(DetectedAttack)++$;

Table A.4: Reward functions for performance variables in meter SAN model

Performance Variable	*Reward Function*
TotalAttacks	$return\ MARK(SharedAttack)$;
UndetectedAttacks	$return\ MARK(UndetectedAttack)$;
DetectedAttacks	$return\ MARK(DetectedAttack)$;
SensorWorking	$return\ MARK(WorkingSensors)$;

Table A.5: Experiment variable values for basic meter model

Experiment Variable	*Value*
Attack Count	1
Coverage Rate	p_c (see Section 3.6.2 and Table 3.1)
DetectAttackRate	950.0
MissAttackRate	50.0
NoCoverageRate	$10 - p_c$
SensorFailureRate	0.05
SensorRepairRate	0.98

Table A.6: Initial markings for places in expanded meter SAN model

Place Name	*Initial Marking*
Attack	0
DetectedAttack	0
ID	0
SharedAttack	0
SharedID	0
SharedWorking	2
UndetectedAttack	0
WorkingSensor	0

Table A.7: Enabling predicates for input gates in expanded meter SAN model

Input Gate	*Predicate*
IG1	$return\ MARK(ID);$
IG_AutoMissAttack	$return\ (!MARK(WorkingSensor)\ \&\&\ MARK(Attack));$
IG_DetectAttack	$return\ (MARK(WorkingSensor)\ \&\&\ MARK(Attack));$
IG_ID	$return\ (!MARK(ID)\ \&\&\ MARK(SharedID));$

Table A.8: Completion functions for output gates in expanded meter SAN model

Output Gate	*Function*
OG1	$intcoverage[16][21] =$ {two-dimensional array of coverage vectors} $if(MARK(SharedWorking) > 0\ \&\&\ coverage[11][MARK(ID)])$ $\quad MARK(WorkingSensor) = 1;$ $else$ $\quad MARK(WorkingSensor) = 0;$
OG_AutoMissAttack	$MARK(Attack)--;$ $MARK(UndetectedAttack)++;$
OG_DetectAttack	$MARK(Attack)--;$ $MARK(DetectedAttack)++;$
OG_ID	$MARK(ID) = MARK(SharedID);$ $MARK(SharedID) = 0;$
OG_MissAttack	$MARK(Attack)--;$ $MARK(UndetectedAttack)++;$

Table A.9: Initial markings for places in expanded meter sensor system SAN model

Place Name	*Initial Marking*
Broken	0
Change	1
ID	0
SharedID	0
SharedWorking	2
Working	1

Table A.10: Enabling predicates for input gates in expanded meter sensor system SAN model

Input Gate	*Predicate*
IG1	$return\ MARK(ID)\ \&\&\ MARK(Change);$
IG_ID	$return\ (!MARK(ID)\ \&\&\ MARK(SharedID));$

Table A.11: Completion functions for output gates in expanded meter sensor system SAN model

Output Gate	*Function*
Change1	$MARK(Change)++;$ $MARK(Broken)++;$
Change2	$MARK(Change)++;$ $MARK(Working)++;$
OG_ID	$MARK(ID) = MARK(SharedID);$ $MARK(SharedID) = 0;$
Update	$MARK(Change) = 0;$ $if(MARK(Working))$ $\quad MARK(SharedWorking) \parallel = MARK(ID);$ $else$ $\quad MARK(SharedWorking)\ \&= \sim MARK(ID);$

REFERENCES

[1] U.S. Energy Information Administration, "What is the electric power grid, and what are some challenged it faces?" http://www.eia.gov/energy_in_brief/article/power_grid.cfm, Apr. 2012.

[2] B. Krebs, "FBI: Smart meter hacks likely to spread," http://krebsonsecurity.com/2012/04/fbi-smart-meter-hacks-likely-to-spread/, Apr. 2012.

[3] S. McLaughlin, D. Podkuiko, and P. McDaniel, "Energy Theft in the Advanced Metering Infrastructure," in *Proceedings of 4th International Conference on Critical Information Infrastructures Security*, 2009.

[4] X. Wu and N. Li, "Achieving privacy in mesh networks," in *Proceedings of the Fourth ACM Workshop on Security of Ad Hoc and Sensor Networks*, 2006, pp. 13–22.

[5] D. Welch and S. Lathrop, "Wireless security threat taxonomy," in *Proceedings of the IEEE Systems, Man and Cybernetics Society Information Assurance Workshop*, June 2003, pp. 76–83.

[6] A. K. Rai, R. R. Tewari, and S. K. Upadhyay, "Different types of attacks on integrated MANET-Internet communication," *International Journal of Computer Science and Security*, vol. 4, no. 3, pp. 265–274.

[7] B. Kannhavong, H. Nakayama, Y. Nemoto, N. Kato, and A. Jamalipour, "A survey of routing attacks in mobile ad hoc networks," *IEEE Wireless Communications*, vol. 14, no. 5, pp. 85–91, Oct. 2007.

[8] A. A. Cardenas, T. Roosta, and S. Sastry, "Rethinking security properties, threat models, and the design space in sensor networks: A case study in SCADA systems," *Ad Hoc Networks*, vol. 7, no. 8, pp. 1434–1447, 2009.

[9] T. Gamer, L. Voelker, and M. Zitterbart, "Differentiated security in wireless mesh networks," *Security and Communication Networks*, vol. 4, no. 3, pp. 257–266, Mar. 2011.

[10] H. Redwan and K.-H. Kim, "Survey of security requirements, attacks and network integration in wireless mesh networks," in *Proceedings of the New Technologies, Mobility and Security, 2008. NTMS '08*, Nov. 2008, pp. 1–5.

[11] X. Wang and P. Yi, "Security framework for wireless communications in smart distribution grid," *IEEE Transactions on Smart Grid*, vol. 2, no. 4, pp. 809–818, Dec. 2011.

[12] S. McLaughlin, D. Podkuiko, and P. McDaniel, "Energy theft in the Advanced Metering Infrastructure," in *Proceedings of the 4th international conference on Critical information infrastructures security*, ser. CRITIS'09. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 176–187.

[13] S. McLaughlin, D. Podkuiko, S. Miadzvezhanka, A. Delozier, and P. McDaniel, "Multi-vendor penetration testing in the Advanced Metering Infrastructure," in *Proceedings of the 26th Annual Computer Security Applications Conference*, 2010, pp. 107–116.

[14] D. Grochocki, J. H. Huh, R. Berthier, R. B. Bobba, W. H. Sanders, A. A. Cardenas, and J. G. Jetcheva, "AMI threats, intrusion detection requirements and deployment recommendations," in *Proceedings of the Third IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Nov. 2012.

[15] D. Boyle and T. Newe, "Security protocols for use with wireless sensor networks: A survey of security architectures," in *Third International Conference on Wireless and Mobile Communications, ICWMC '07*, Mar. 2007, p. 54.

[16] A. Perrig, R. Szewczyk, J. D. Tygar, V. Wen, and D. E. Culler, "SPINS: Security protocols for sensor networks," *Wireless Networks*, vol. 8, no. 5, pp. 521–534, Sept. 2002.

[17] S. Zhu, S. Setia, and S. Jajodia, "Leap: efficient security mechanisms for large-scale distributed sensor networks," in *Proceedings of the 10th ACM Conference on Computer and Communications Security*, ser. CCS '03. New York, NY, USA: ACM, 2003, pp. 62–72.

[18] G. Dini and M. Tiloca, "Considerations on Security in ZigBee Networks," in *Proceedings of the IEEE International Conference on Sensor Networks, Ubiquitous, and Trustworthy Computing (SUTC)*, June 2010, pp. 58–65.

[19] M. Fouda, Z. Fadlullah, and N. Kato, "Assessing attack threat against ZigBee-based home area network for smart grid communications," in *Proceedings of the International Conference on Computer Engineering and Systems (ICCES)*, Dec. 2010, pp. 245–250.

[20] S. Zhu, S. Setia, and S. Jajodia, "LEAP+: Efficient security mechanisms for large-scale distributed sensor networks," *ACM Transactions on Sensor Networks*, vol. 2, no. 4, pp. 500–528, Nov. 2006.

[21] C. Gomez and J. Paradells, "Wireless home automation networks: A survey of architectures and technologies," *IEEE Communications Magazine*, vol. 48, no. 6, pp. 92–101, June 2010.

[22] P. Shamsi and B. Fahimi, "Remote control of smart appliances using MPEI," in *Power Engineering, Energy and Electrical Drives (POWERENG), 2011 International Conference on*, May 2011, pp. 1–5.

[23] S. Wolff, P. G. Larsen, K. Lausdahl, A. Ribeiro, and T. S. Toftegaard, "Facilitating home automation through wireless protocol interoperability," in *Proceedings of the 12th International Symposium on Wireless Personal Multimedia Communications (WPMC)*, 2009.

[24] F. Siddiqui, S. Zeadally, C. Alcaraz, and S. Galvao, "Smart grid privacy: Issues and solutions," in *Proceedings of the 21st International Conference on Computer Communications and Networks (ICCCN)*, July 2012, pp. 1–5.

[25] P. McDaniel and S. McLaughlin, "Security and privacy challenges in the smart grid," *IEEE Security Privacy*, vol. 7, no. 3, pp. 75–77, May-June 2009.

[26] Y. Simmhan, A. Kumbhare, B. Cao, and V. Prasanna, "An analysis of security and privacy issues in smart grid software architectures on clouds," in *Proceedings of the IEEE International Conference on Cloud Computing (CLOUD)*, July 2011, pp. 582–589.

[27] R. Chow, E. Uzun, A. A. Cardenas, Z. Song, and S. Lee, "Enhancing cyber-physical security through data patterns," in *Proceedings of the Workshop on Foundations of Dependable and Secure Cyber-Physical Systems*, Apr. 2011.

[28] B. Vaidya, D. Makrakis, and H. Mouftah, "Device authentication mechanism for smart energy home area networks," in *Proceedings of the IEEE International Conference on Consumer Electronics (ICCE)*, Jan. 2011, pp. 787–788.

[29] R. Berthier, W. H. Sanders, and H. Khurana, "Intrusion detection for Advanced Metering Infrastructures: Requirements and architectural directions," in *Proceedings of teh 2010 First IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Oct. 2010, pp. 350–355.

[30] P. Jokar, H. Nicanfar, and V. Leung, "Specification-based intrusion detection for home area networks in smart grids," in *Proceedings of the 2011 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Oct. 2011, pp. 208–213.

[31] D. Jin, C. Lee, D. Nicol, I. Shin, and H. Zhu, "Simulation-based study of distributed denial-of-service attacks in Advanced Metering Infrastructure," in *INFORMS Annual Meeting*, Charlotte, NC, USA, Nov. 2011.

[32] I. Shin, J. H. Huh, C. Lee, and D. M. Nicol, "A Monitoring Architecture for Smart Meter Mesh Networks in the Smart Grid," 2011.

[33] D. D. Deavours, G. Clark, T. Courtney, D. Daly, S. Derisavi, J. Doyle, W. Sanders, and P. Webster, "The Möbius framework and its implementation," *IEEE Transactions on Software Engineering*, vol. 28, no. 10, pp. 956–969, Oct 2002.

[34] "U.S. Census 2000 Summary File 1 - Population, Housing Units, Area, and Density: 2000."

[35] Novarum, Inc., "2010 guidelines for successful large scale outdoor Wi-Fi networks," http://www.scribd.com/doc/25026923/2010-Guidelines-for-Large-Scale-Outdoor-WiFi, Dec. 2009.

[36] *ANSI C12. Smart Grid Meter Package*, ANSI Std., 2011.

[37] A. W. Moore and D. Zuev, "Internet traffic classification using Bayesian analysis techniques," in *ACM SIGMETRICS Performance Evaluation Review*, vol. 33, no. 1, 2005, pp. 50–60.

[38] A. De Montigny-Leboeuf, "Flow attributes for use in traffic characterization," *Communications Research Centre Canada, Tech. Rep*, 2005.

[39] T. Karagiannis, K. Papagiannaki, and M. Faloutsos, "BLINC: Multilevel traffic classification in the dark," in *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 4, 2005, pp. 229–240.

[40] L. Bernaille, R. Teixeira, I. Akodkenou, A. Soule, and K. Salamatian, "Traffic classification on the fly," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 2, pp. 23–26, 2006.

[41] L. Bernaille, R. Teixeira, and K. Salamatian, "Early application identification," in *Proceedings of the 2006 ACM Conference on emerging Networking EXperiments and Technologies*, 2006, p. 6.

[42] L. Bernaille and R. Teixeira, "Early recognition of encrypted applications," in *Passive and Active Network Measurement*. Springer, 2007, pp. 165–175.

[43] S. Roschke, L. Ibraimi, F. Cheng, and C. Meinel, "Secure communication using identity based encryption," in *Communications and Multimedia Security*. Springer, 2010, pp. 256–267.

[44] N. Williams, S. Zander, and G. Armitage, "A preliminary performance comparison of five machine learning algorithms for practical IP traffic flow classification," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 5, pp. 5–16, 2006.

[45] M. Crotti, M. Dusi, F. Gringoli, and L. Salgarelli, "Traffic classification through simple statistical fingerprinting," *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 1, pp. 5–16, 2007.

[46] G.-L. Sun, Y. Xue, Y. Dong, D. Wang, and C. Li, "A novel hybrid method for effectively classifying encrypted traffic," in *Proceedings of the IEE Global Telecommunications Conference (GLOBECOM)*, 2010, pp. 1–5.

[47] M. Kharrazi, S. Sen, and O. Spatscheck, "Towards real-time performance monitoring for encrypted traffic," in *Proceedings of the 2007 SIGCOMM Workshop on Internet Network Management*, 2007, pp. 287–292.

[48] V. T. Goh, J. Zimmermann, and M. Looi, "Towards intrusion detection for encrypted networks," in *Proceedings of the International Conference on Availability, Reliability and Security, 2009. ARES'09*, 2009, pp. 540–545.

[49] V. T. Goh, J. Zimmermann, and M. Looi, "Intrusion detection system for encrypted networks using secret-sharing schemes," *International Journal of Cryptology Research*, 2010.

[50] A. F. Esteves, P. R. Inácio, M. Pereira, and M. M. Freire, "On-line detection of encrypted traffic generated by mesh-based peer-to-peer live streaming applications: The case of GoalBit," in *Proceedings of the 10th IEEE International Symposium on Network Computing and Applications (NCA)*, 2011, pp. 223–228.

[51] R. Koch and G. D. Rodosek, "Command evaluation in encrypted remote sessions," in *Proceedings of the 4th International Conference on Network and System Security (NSS)*. IEEE, 2010, pp. 299–305.