SANDIA REPORT

SAND2019-2406
Unlimited Release
Printed February 2019

Cybersecurity Assessments on Emulated DER Communication Networks

Ifeoma Onunkwo, Patricia Cordeiro, Brian Wright, Nicholas Jacobs, Christine Lai, Jay Johnson, Trevor Hutchins, William Stout, Adrian Chavez, Bryan T. Richardson, Keith Schwalm

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology and Engineering Solutions of Sandia, LLC.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.osti.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online ordering: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2019-2406
Unlimited Release
Printed February 2019

Cybersecurity Assessments on Emulated DER Communication Networks

Ifeoma Onunkwo & Brian Wright
Computer Systems Security Analysis R&D
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0932
{ionunkw, bjwrigh}@sandia.gov

Patricia Cordeiro
Analytics and Cryptography
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-1027
pgcorde@sandia.gov

Nicholas Jacobs & Christine Lai
Cyber Resilience R&D
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0757
{njacobs, cflai}@sandia.gov

Jay Johnson
Renewable& Distributed Systems Integration
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-1033
jjohns2@sandia.gov

Trevor Hutchins
Critical Infrastructure Systems
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0671
thutchi@sandia.gov

William Stout
Cyber Security Initiatives
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0813
wmstout@sandia.gov

Adrian Chavez
Autonomous Cyber Systems
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0672
adrchav@sandia.gov

Bryan T. Richardson & Keith Schwalm
DNK Consulting
12300 Crested Moss Road NE
Albuquerque, NM 87112
keith@dnk.com & bryan@activeshadow.com

Abstract

An increasing number of public utility commissions are adopting Distributed Energy Resource (DER) interconnection standards which require photovoltaic (PV) inverters, energy storage systems, and other DER to include interoperable grid-support functionality. The recently updated national standard, IEEE 1547-2018, requires all DER to include a SunSpec Modbus, IEEE 2030.5, or IEEE 1815 communication interface in order to provide local and bulk power system services. Those communication protocols and associated information models will ensure system interoperability for PV and storage systems, but these new utility-to-DER communication networks must be deployed with sufficient cybersecurity to protect the U.S. power system and other critical infrastructure reliant on dependable power. Unlike bulk generators, DER are commonly connected to grid operators via public internet channels. These DER networks are exposed to a large attack surface that may leverage sophisticated techniques and infrastructure developed on IT systems, including remote exploits and distributed attacks. Although DER make up a growing portion of the national generation mix, they have limited processing capabilities and do not typically support modern security features such as encryption or authentication.

In this work, Sandia National Laboratories constructed simulated DER communication networks with a range of security features in order to study the security posture of different communication approaches. The experimental test environment was created in a Sandia-developed co-simulation platform, called SCEPTRE, which emulated SunSpec-compliant DER equipment, the utility DER management system, communication network, and distribution power system. Adversary-based assessments were conducted and a quantitative scoring criteria was applied to evaluate the resilience of various architectures against cyber attacks and to measure the systemic impact during such attacks. The team found that network segmentation, encryption, and moving target defense improved the security of these networks and would be recommended for utility, aggregator, and local DER networks.

Acknowledgment

The authors would like to thank the SCEPTRE development team for support during the project, and making a number of updates to the source-code to complete the assessments. We would especially like to thank Matthew Reno for helping with the OpenDSS integration, and Derek Hart and Jordan Henry for the SunSpec RTU updates.

Contents

List of Figures

List of Tables

Nomenclature

API Application Program Interface

DER Distributed Energy Resource

Emulytics™ A trademarked term referring to the combination of Emulation and Analytics within a set of capabilities developed by Sandia National Laboratories

FEP Front End Processor

ICS Industrial Control System

OPC Open Platform Communications

PLC Programmable Logic Controller

PV Photovoltaic, such as a Photovoltaic Inverter

RTU Remote Terminal Unit

SCADA Supervisory Control And Data Acquisition

SCEPTRE A comprehensive ICS/SCADA modeling and simulation platform that captures the cyber/physical impacts of targeted cyber events on critical infrastructure and control systems

TLS Transport Layer Security

This page intentionally left blank.

Chapter 1

Introduction

There is ample evidence from the last few years that many power system networks in the US [1, 2, 3, 4, 5, 6] and abroad [7] are the target of active cybersecurity reconnaissance and attacks. The most widely publicized attacks are those that caused widespread blackouts in Ukraine in 2015 and 2016 [8, 9], but there have been several other disconcerting trends including: the increase in operation technology (OT)-focused malware, e.g., Crash Override and Black Energy [10, 11], deep reconnaissance into power system networks [12, 13, 14], and growing willingness to deploy powerful cyber weapons that are affecting critical infrastructure [8, 9, 15]. Attackers often use myriad techniques to gain footholds in information technology (IT) networks and then pivot to other computers, servers, and networks to exfiltrate sensitive information, monitor operations, or plan for sophisticated attacks [16].

At the same time, penetrations of Distributed Energy Resources (DER), e.g., Photovoltaics (PV) and Energy Storage Systems (ESS), continue to grow rapidly on distribution and subtransmission systems [17, 18]. Over the last decade, an increasing number of inverter vendors and aggregators have provided monitoring portals for their customers. Like many other Internet of Things (IoT) devices, customers can monitor or control their equipment via proprietary communication protocols. However, this IoT equipment now controls a substantial portion of the total power production in certain jurisdictions, like Hawaii and California [19, 20].

In 2018, a revision to the US interconnection and interoperability standard, IEEE Std. 1547, required DER equipment to have either an IEEE 2030.5 (SEP 2), IEEE 1815 (DNP3), or SunSpec Modbus communication interface [21]. New California Public Utility Commission (CPUC) Electric Rule 21 regulations that go into effect in early 2019 define IEEE 2030.5 [22] as the default application protocol for Investor Owned Utilities (IOUs) communications to DER [23, 24]. The adoption of standardized communication protocols is a critical step toward interoperability between power system operators and DER equipment, but a comprehensive national approach to DER cybersecurity is absent. At this point, the network architecture and cybersecurity requirements are not fully defined for California, and other states will also have to make similar decisions as utilities and other grid operators start interacting with interoperable DER. In this report, the team researched communication requirements for different grid applications in order to generate reference architectures and cybersecurity recommendations for discussions in California and elsewhere regarding the best methods to take advantage of increasing PV penetrations and advanced inverter capabilities.

There are many security requirements for operators of critical infrastructure in the US. Power system operators are required to adhere to the North American Electric Reliability Corporation (NERC) Critical Infrastructure Protection (CIP) standards which cover—among other things—training, security and information management, perimeter defenses, and incident reporting [25]. NREC requirements are reserved for bulk power equipment operating at or above 100 kV, so DER equipment and associated networks are exempt from these requirements. The solar industry and national government understand this gap in power system security and are working to address the requirements by reviewing and updating security requirements in the DER communication protocols [26, 27], standing up DER cybersecurity working groups [28], and seeking new security standards for DER devices and networks [29].

There are a wide range of R&D areas that may improve the national DER cybersecurity posture [30]. In this work, three network defense techniques were analyzed with respect to confidentiality, integrity, and availability (CIA) - the three tenets of cybersecurity. Network segmentation, encryption, and moving target defense (MTD) were deployed in a virtualized environment to quantify their security improvements in the broad CIA areas by conducting adversary-based (red team) assessments. A calculation of additional latency associated with these features was also conducted to determine if these security features would interfere with grid operations supported by DER. This work produced power system cybersecurity metrics to advise the solar and power system industry on best cybersecurity practices for DER networks.

Chapter 2

Emulytics Simulation Environment

To demonstrate and evaluate the cybersecurity implications of communication for DER support of the grid, an Emulytics^{TM} environment was created and utilized to combine power simulation with virtual networking using the SCEPTRE platform developed by Sandia National Laboratories. In describing this environment, first the underlying power model and simulation will be described in Section 2.1. Then the components of the SCEPTRE platform utilized in this research will be discussed in Section 2.2.

2.1 Power Simulations

To examine the cybersecurity implications for communications within a control network and its impact on the complete system, analysis of the networking is insufficient. Rather, a combination of the networking and an underlying power system simulation is required. This section will discuss two power system models that were created for the red team assessments.

The first model represented a 12 kV distribution feeder with two 750 kW PV sites, as shown in Figure 2.1. This model was anonomyzed from a utility partner and developed for voltage regulation studies [31, 32]. The feeder included 215 buses and 39 transformers, and had a peak load of 3.98 MW. The GridPV toolbox in OpenDSS [33] was used to conduct quasi-static time series (QSTS) simulations for an October day with minimum daytime load (1.51 MW) to maximize the instantaneous penetration of PV. The original OpenDSS model represented the two PV sites with single 750 kW inverters, but these were each replaced with 10 inverters with 75 kVA nameplate capacities. This keep the aggregated generation capacity the same, but provided 20 individually controllable and addressable inverters to be networked into different ways to so the benefits of different security architectures.

A transmission model was also created using the 42-bus transmission model with 26 PV systems, as shown in Figure 2.2. While this model was implemented in SCEPTRE as a proof-of-concept for transmission co-simulations with high penetration PV scenarios, it was not used in any of the red team assessments.

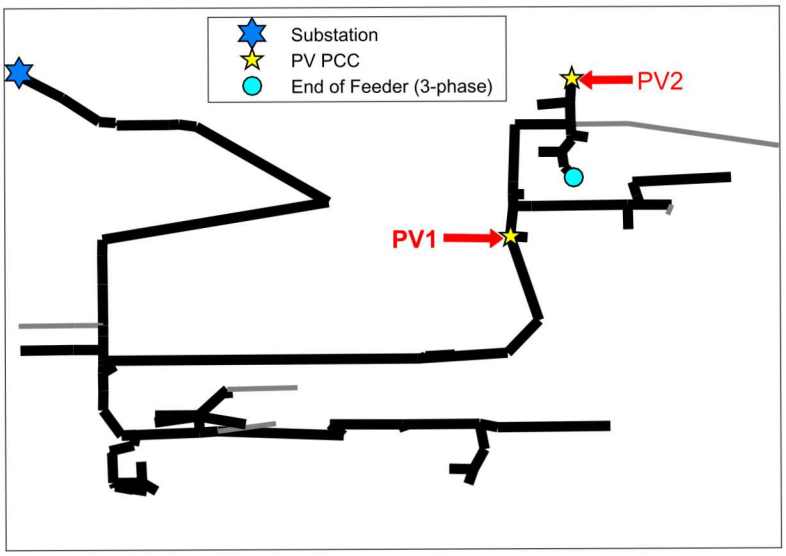

Figure 2.1. 12 kV distribution feeder used in the the SCEPTRE red team experiments.

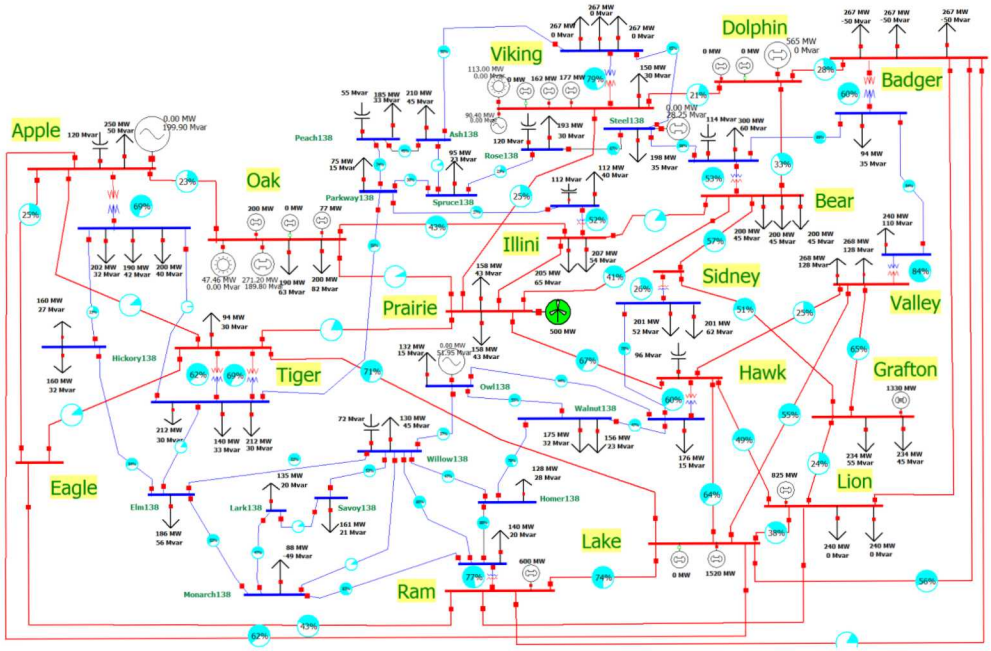

Figure 2.2. Transmission PowerWorld model for SCEPTRE experiments.

2.2 SCEPTRE

SCEPTRE is an Emulytics™ platform developed by Sandia National Laboratories for co-simulation of Industrial Control Systems with virtualized networking. It has been developed to allow for variable fidelity environments, using system simulations, emulated devices, and Hardware-In-The-Loop (HIL). To accomplish this, SCEPTRE uses an underlying simulation model which can utilize many known and common simulation platforms (such as OpenDSS or PowerWorld). These simulations are then connected back into a virtual network representing the corresponding control system. This virtual network is created by another Sandia tool called Minimega [34] which has been fully integrated with SCEPTRE. SCEPTRE then takes these two systems and transparently interconnects them so that the network emulation and underlying system simulation may interact just as a control system would interact with the underlying physical system it controls. This interaction is generally completed using virtual machines that represent Remote Terminal Units (RTU), and are thus named such in the environment.

In Figure 2.3, three domains are shown which represent the SCEPTRE network emulation (top), Minimega-control virtual machines and RTUs (middle), power system simulation provider (bottom). In the upper layer, the network directs actual communication packets through emulated routers, switches, and other networking components. In the middle layer, devices in this network (utility DERMS, DER, etc.) are built on Windows and Linux VMs. The power simulation at the bottom, updates values represented in the VMs/RTUs in the middle layer, and also updates the power simulation based on changes to those devices through an unexposed back-end network. On the right, the HIL DER is shown. For the HIL SCEPTRE experiment in this project, the PV inverter was connected to a real PV array but the AC power was provided with a grid simulator. As changes in the power system occurred in the provider (e.g., a new voltage level), those changes were communicated to the Ametek RS180 grid simulator to change the terminal voltage on the inverter. The active and reactive power of the inverter were calculated using a LabVIEW data acquisition system (DAS) connected to current and voltage probes. From those measurements, the power factor of the inverter was calculated and updated in the OpenDSS simulation.

The rest of this section will discuss these components in greater detail. Section 2.2.1 will discuss the use of a power simulation within SCEPTRE and Section 2.2.2 will discuss the RTUs representing DERs in our models. The implementation of encryption and other potential security measures within SCEPTRE are presented in Sections 2.2.3 and 2.2.3. Finally, Section 2.2.4 will present the various control network topologies represented and utilized for the Red Teaming activities of Chapter 4.

2.2.1 Power Simulation Interaction

The interface between a power simulation and SCEPTRE is done through what is known as a "Provider"in SCEPTRE. This code uses any Application Program Interface (API) available, or develops one, to connect to a simulation platform such as OpenDSS. The provider then connects back into SCEPTRE through a standard communication interface with a Publish/Subscribe model

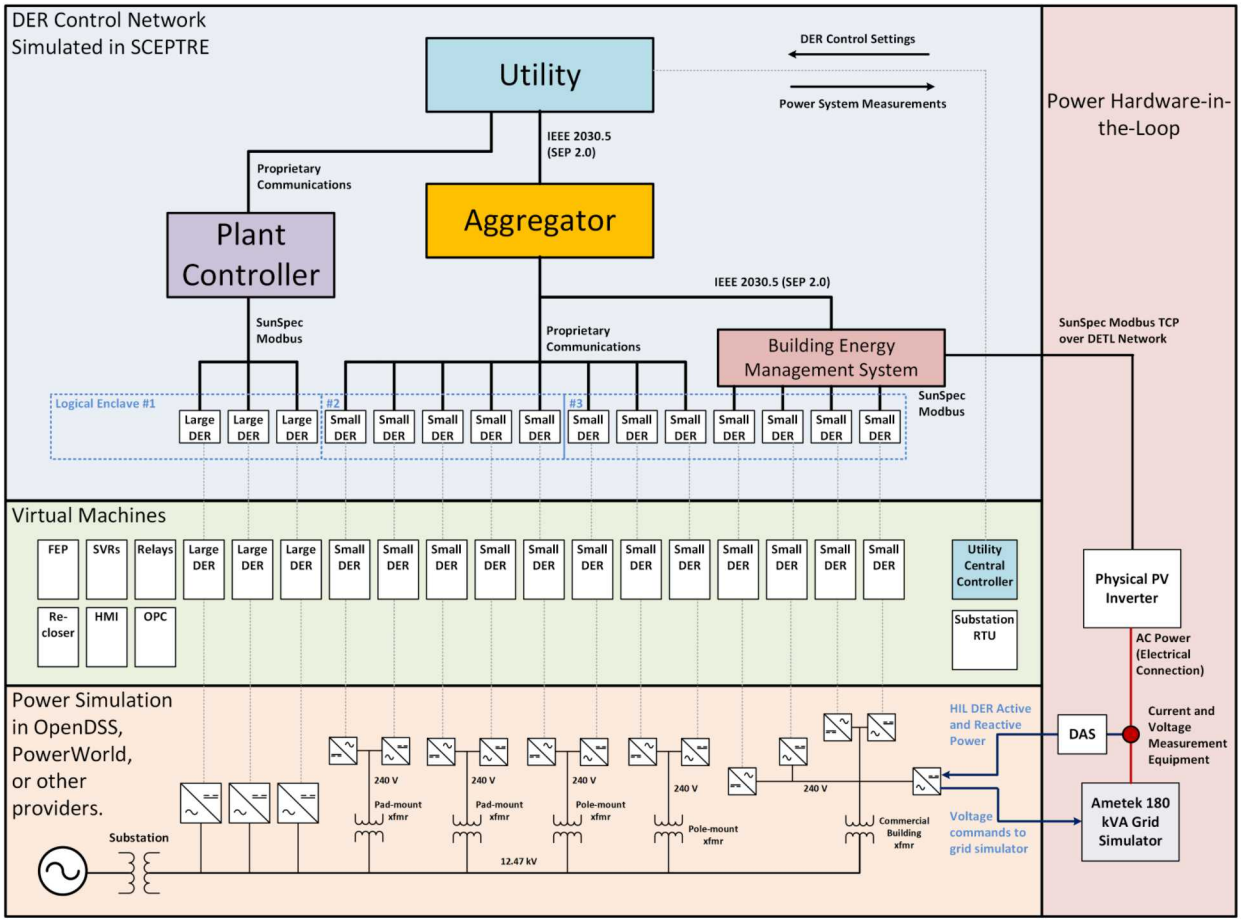

Figure 2.3. Overview of the SCEPTRE environment with a HIL DER.

using ZeroMQ.

In the case of the power hardware-in-the-loop (PHIL) experiment, a physical PV inverter was connected to the power simulation as shown in Fig. 2.3. The OpenDSS simulation was run and power system parameters (voltage, frequency, DER current, etc.) were passed through ZeroMQ to the virtualized devices to update their internal Modbus maps with the latest data. In the case of the HIL inverter, the voltage was sent to the Ametek RS 180 grid simulator over a TCP/IP connection using IEEE-488.2 (SCPI) commands to update the AC voltage applied to the inverter. A Windows computer running a LabVIEW program connected to current and voltage probes on the AC terminals of the PV inverter and acted as the Data Acquisition System (DAS). The LabVIEW program broadcast UDP packets containing active and reactive power measurements to the SCEPTRE environment every second. This data was captured by SCEPTRE and used to update the active and reactive powers of the PV inverter in the OpenDSS power simulation with measurements from the physical inverter. Therefore, the OpenDSS power simulation was resolved every second with new data from the physical inverter. Interestingly, the DAS measurements needed to be used as opposed to those in the Ametek, because the Ametek measurements of active and reactive power were not accurate–likely because the probes were not located at the terminals of the inverter and this was a small PV inverter compared to the grid simulator capacity.

2.2.2 Remote Terminal Units

Within SCEPTRE, the RTUs are virtual machines (VMs) that perform the following functions.

1. Represent field devices in an ICS or SCADA network
2. Connect and interact transparently with the underlying system simulation
3. Perform simple control logic
4. Communicate with the control network using a standard protocol, such as Modbus/TCP or DNP3

To represent a field device in a SCADA network, SCEPTRE RTUs are generally operated as pared down versions of Linux with select capabilities. This better represents the limited functionality of a field device that is only completing a few specific tasks. Furthermore, it allows for an increased number of RTUs to be created within a SCEPTRE experiment without requiring more hardware resources to be allocated. The RTU communicates with the power simulation using a ZeroMQ messaging framework on the backend of SCEPTRE. This process uses a Publish/Subscribe model to both pull data from the power simulation and push control settings and commands back into the simulation. RTUs are programmed with simple control logic representative of its role in the emulated SCADA network. This control logic and processing should be equivalent to that which would be on a field device in the corresponding real-world system.

The RTU, as well as any other components in the ICS network such as Programmable Logic Controllers (PLC), Front End Processors (FEP), or an Open Platform Communications (OPC)

Server, can communicate commonly used control protocols on the ICS network such as Modbus/TCP or DNP3. This means that the communications used within the control network should be represented and utilized to communicate information to the appropriate locations. These communications and network connections are analyzed in the Red Teaming activities of Chapter 4.

Measurement and grid-support functionality was added to SCEPTRE RTUs to represent the capabilities of a SunSpec-compliant smart inverters using SunSpec Information Models 1, 101, 123, and 126, which represent the Common, Inverter (Single Phase), Immediate Controls, and Static Volt-VAR data [35]. The associated connections between the inverter and OpenDSS and PowerWorld were constructed across the ZeroMQ network.

2.2.3 Security Mechanisms

To perform cybersecurity analyses of the impact of the security control mechanisms on both the security of the DER control system and the corresponding impact to the grid, the implementation of various security measures was required. This section discusses those mechanisms examined in this work and their implementation within the emulated control network.

Network Segmentation and Enclaving

Network segmentation is a known and commonly used strategy for providing additional security to a control network. The extreme example of such a control mechanism is the "Air Gap", where a control network is isolated by physically separation from any other network. In reality, maintenance of a control network commonly requires crossing such an air gap. To address that need, control networks may be logically or physically segmented to reduce the impact of a security compromise to a subset of the system. This can be done through physical segmentation of the control network, or through mechanisms such as Virtual Local Area Networks (VLANs).

Network segmentation is a technique to minimize common-mode vulnerabilities, whereby enclaves are isolated with firewall rules, VPNs, proxies, or other networking technologies so that traffic between them is only allowed by exception. Extensive research on segmentation for military microgrids has been completed previously [36, 37], and the enclaving concepts represented here were derived from [36, 38], but have been modified for an ADMS application. The downside of this approach is the additional network administration and communication latency. There are technical challenges to segment DER networks because networking equipment will not necessarily be owned and operated by a single entity. It may be possible to enclave the devices if communications are passed directly to the DER through networks that are owned by the grid operator, e.g., through an advanced metering infrastructure (AMI) mesh radio or dedicated SCADA network to DER systems. However, in the majority of commercial and residential PV systems, communications will be established through wired or wireless networks via the public internet, as shown in Fig. 2.4. In those cases, it is more difficult to segment or enclave the networks because internet service providers (ISPs) control the network routing and firewall rules, and cannot be implemented

easily without assistance from the ISPs. Therefore, the use of VPNs, proxies, or some other technology would be required to logically isolate the enclaves. Three options and their pros and cons are presented in 2.1. These include using firewall rules—which will only work when operating over a private network, using hardware proxies to hide traffic, and using encrypted VPN tunnels.

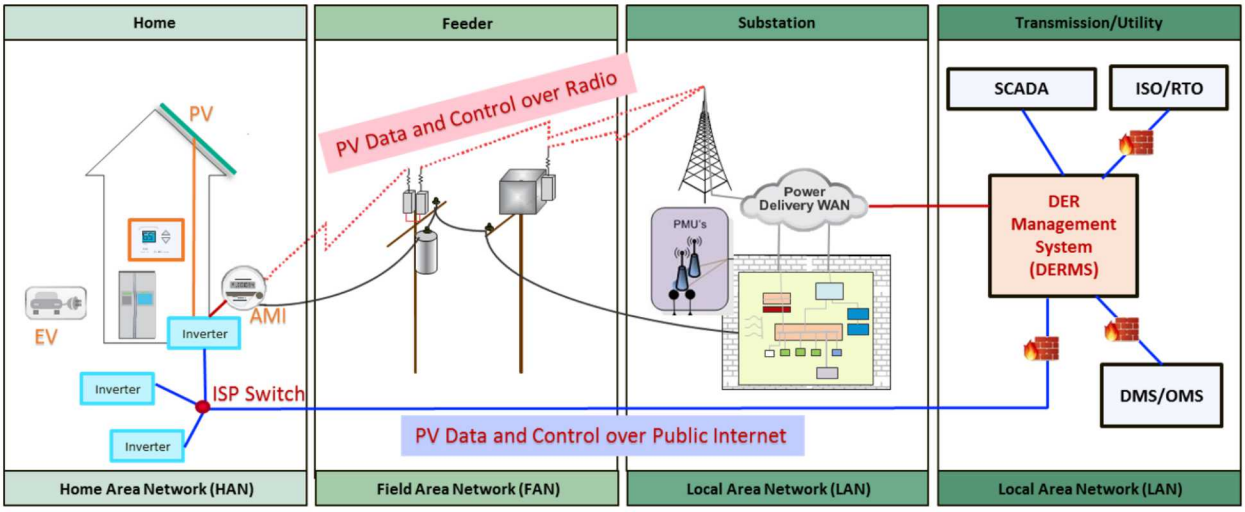

Figure 2.4. Different DER control network architectures in which DER data is exchanged over public internet or AMI radio networks.

Firewall rules have been used in the past for military systems [37], but these are not effective when operating with internet connected devices because the network traffic channels are not consistent and ISP systems are designed for speed, not security. This method becomes an option when the utility or other grid operator is communicating to DER equipment through a network that they own. In that case, they may apply specific firewall rules to create enclaves. Blocking all connections initially and then allowing specific ones is considered a best practice with firewall rules. When the network is privately controlled, this approach will allow a utility to whitelist traffic from a DERMS to each respective HAN (or commercial/utility DER LAN or enclave), and all other traffic would be dropped. This can be easy to manage if the number of HANs is small; however, if that number is extensive and continually changing, it would become a difficult operational management issue. Constantly changing firewall rules can also introduce a greater chance of dropping of legitimate connections by mistake thereby causing significant ongoing support concerns. The advantage—almost necessity—of using a firewall on a private network is to ensure that only traffic desired from the DERMS to a specific HAN can transit the network, and all other traffic (potentially malicious) is dropped.

Another option would be for the utility to provide each physical site with a hardware proxy between the ISP connection and HAN or facility LAN. A hardware proxy is simply a small device like a cable/DSL modem that would have two primary connection. One connection would be to receive the general ISP connection, and the other would be to output to the HAN. Additional connections would be required if the device were intended to connect directly to an inverter—in

Table 2.1. Methods of enclaving DER networks.

Enclave mechanism	Pros	Cons
Firewall rules (whitelist DER-to-headend connections) for grid operator or aggregator-owned networks (e.g., AMI networks)	• Private network • Extends grid operator or aggregator LAN to FAN or HAN	• More costly • More management • Complex • Potentially less data bandwidth or speed
Use hardware proxy, which monitors for DER/utility traffic and exchanges it	• Works well for aggregators • Traffic send via ISPs using RESTful HTTP or TLS connections	• Relies on 3rd party (ISP) to manage network (could have more latency if QoS is an issue) • Need maintenance contracts • Privacy concerns (for unencrypted traffic) • Less flexibility
Virtual private networks (VPNs) between DER and grid operator	• Direct connections between DER and utility • Reduced latency • Grid operator controls and easily changes segmentation	• VPN management and maintenance difficult • Could burden facility/home bandwidth

that case, the hardware proxy would route traffic identified as intended for the inverter before passing it off to the HAN. This proxy would monitor for traffic specific to the inverter and pass that traffic directly to it; all other traffic would be passed unmonitored to the HAN. Controlling specific traffic between the DERMS and an individual (or group) HAN would be similar to the firewall option over a private utility network. There would be potential privacy concerns if the proxy were compromised by an adversary who could then manipulate the network traffic for their benefit. A challenge with the hardware proxy is to install and support a physical hardware device at each site; this could present additional support and maintenance cost to the utility. It may be possible to provide each facility with an ISP-friendly switch/modem in place of what the ISP has provided (most markets have few ISPs to pick from, and so it might be easy to provide an ISP-friendly hardware device). The hardware proxy would also allow for priority traffic specific to the inverters, i.e., priority over regular HAN traffic. Finally, the hardware proxy would need to update the utility through a dynamic-DNS-like service so that the utility was always aware of the (potentially changing) publicly routable IP address of the home or facility.

Alternatively, the utility could maintain an ongoing virtual private network (VPN) connection directly to the inverters through the existing ISP network and corresponding switch/modem. A VPN is an encrypted tunnel for communication between two systems over a network. This would provide the utility with a direct, secure connection between the DERMS and each HAN over a public network based on well-established open standards. Communication encryption prevents eavesdropping or manipulation by an outside party. Traffic specific to each HAN could be communicated through the VPN tunnel with the assurance that it remains secured from any malicious actors along the communication path (similar to traffic sent over a private network). Additional support and maintenance requirements would be necessary from the utility, similar to the hardware proxy, but additional support from a HAN's ISP would not be necessary as most ISPs support VPN tunnels for their customers without any additional service changes. To initiate a new connection, each inverter would initiate the VPN to a known utility IP address providing a "plug-and-play" deployment. An alternative to an inverter initiated connection would be to deploy a facility gateway where the VPN connection could be originated to each respective HAN.

Once the method of generating the enclaves is selected, the specific rules for cleaving the devices must be decided. There must be a balance: too many enclaves mean higher likelihood of mistakes in the firewall, VPN, proxies, or router configurations, slower communication times, and more difficulty deploying more DER; but at the same time, there must be a certain number of enclaves to prevent control of a critical magnitude of generation. Here, we offer two basic approaches:

- A segmented network with DER placed in one of three enclaves at random or convenience, e.g., Fig. 2.5
- A critically segmented network with no more than 20% of the total capacity in a single enclave, e.g., Fig. 2.6

The placement of DER in a specific enclave could be done based on geography, power system topology, nameplate capacity, or other metadata. More sophisticated methods of determining the number of enclaves and which DER should be placed in each should be considered an area of future research. These approaches were initially explored in an earlier Virtual Power Plant project [38].

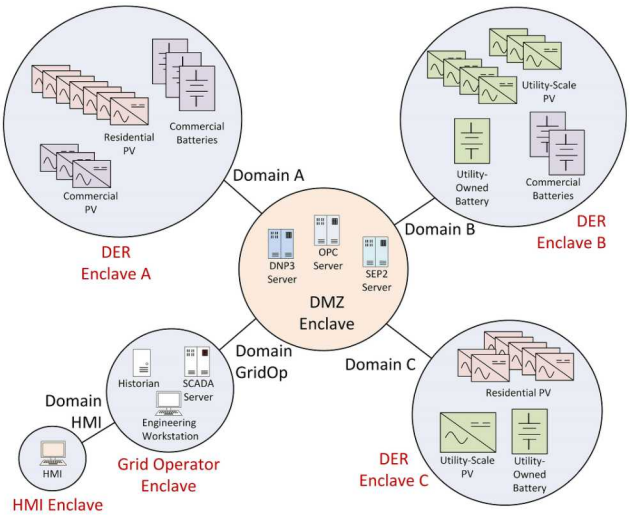

Figure 2.5. Segmented network with DER placed in random enclaves.

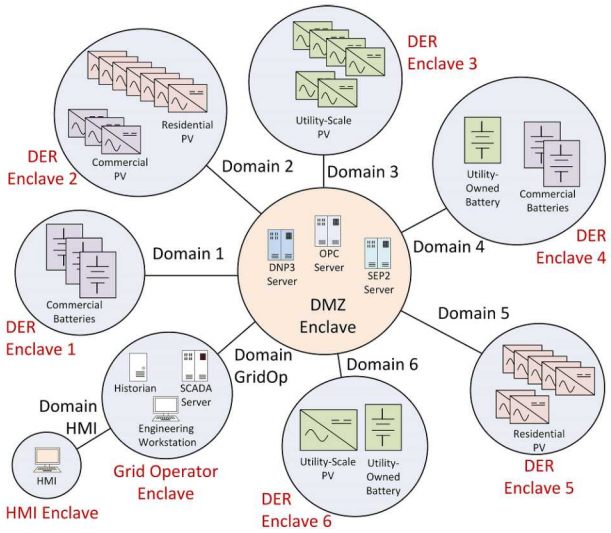

Figure 2.6. Segmented critical network with no more than 20% of the generation placed in one enclave.

Denial of Service (DoS) risk is reduced with network segmentation. Breaking up a network reduces the overall target space by transferring risk from a single network segment to multiple enclaves. As a result, an adversary would need a foothold inside each enclave to launch a DoS attack on the entire DER fleet. This may be relatively trivial for sophisticated adversaries, but it will require additional work on their behalf. Unfortunately, a DoS attack against the utility server as a single target could result in stopping all inbound and outbound DER-related traffic. Using whitelisted firewall rules in the DER infrastructure would be a useful—but time-consuming—solution on the utility side to further reduce the risk from DoS attack related noise. Network segmentation would also reduce the risk of the DER equipment becoming a part of a botnet used for distributed DoS (DDoS) attacks because it increases the effort required to reach the devices. To assess this approach, co-simulations of a distribution power system and DER control network were created and run. Adversary-based assessments were then conducted on these environments to investigate the effectiveness of enclaving techniques to a range of attacks including DoS, as described in 4.

Encryption

As SCEPTRE utilizes known and common control protocols to communicate, it is also capable of integrating common encryption schemes into those communications. The SunSpec-compliant DER inverter RTUs communicate Modbus TCP to the utility ADMS (SVP) VM. Since Modbus is passed in plaintext, it was desired to encrypt these communications to make the environment more realistic. The most typical way to accomplish this, and the method used in this research, is to incorporate Transport Layer Security (TLS) to secure communications using a common implementation such as OpenSSH [39]. Therefore, in some of the environments SSH components were added to the environments to pass the Modbus traffic between the utility and the DER subnets through an encrypted tunnel. This additional security mechanism was performed for the two enclaving strategies.

Moving Target Defense (MTD)

Moving target defense (MTD) is class of technology that dynamically modify a system environment to create uncertainty for adversaries. Chavez et al. developed Artificial Diversity and Dynamic Security (ADDSec) as a MTD tool that leverages software defined networking (SDN) to randomize network parameters and communication paths. ADDSec has the ability to randomize IP addresses and port numbers both in anticipation of and in response to detected network activity. This is meant to thwart an adversary's ability to conduct reconnaissance and establish communications between devices on the network, and has been proven to be effective at increasing the resilience of grid wide area networks against certain types of attacks [40]. ADDSec is comprised of several components:

1. Dynamic reconfiguration of networking and routing parameters, using pseudo-random number generators as a source of entropy, including randomization of IP addresses and ports, to thwart reconnaissance and prevent unwanted connections.

2. Generation of unique application binaries within a system to raise the difficulty of producing software exploits.
3. An ensemble of machine learning algorithms that analyze host statistics, networking information, and network traffic to autonomously detect and trigger reconfiguration of the systems in real-time.

For more information on how ADDSec operates, such as its use of software defined networking (SDN) to enable transparency to end-devices, please refer to [40, 41]. Intuitively, the use of dynamic configurations to decrease predictability for attackers seems reasonable as a means for enhancing cyber security, but techniques to measure the resilience benefits of MTDs to-date has been primarily survey- and opinion-based [42]. Integrating ADDSec in our emulated system topologies allows us to evaluate its effectiveness against a controlled baseline. It is also worth noting that while ADDSec has been proven effective against reconnaissance attacks, it does not necessarily provide protection against a persistent threat that has previously been introduced to the network or prevent an attacker from carrying out an exploit on a previously compromised host.

An example of this technology is shown in Figure 2.7. On the left is a utility subnet consisting of an Advanced Distribution Management System (ADMS), Geographical Information System (GIS), and DER management system (DERMS). On the right, is a collection of DER in a campus or utility/commercial site on a single switch. There is an "IP Generator" computer in the bottom that sends the new IP addresses to the switches in front of actual DER or computation devices. The MTD changes the IP addresses of these switches but the utility-owned and DER nodes retain static IP addresses. Actual implementation would likely require multiple MTD subsystems that independently reconfigure the IP addresses of the utility subnet and DER devices. Since this technology requires a sperate network to be overlaid on the publicly-addressable one, it is likely that residential DER will not have the ability to be included in the MTD/SDN overlay, but commercial and utility owned sites could employ this technology.

2.2.4 Topologies

This section enumerates the different cybersecurity reference architectures used for the red team evaluations. The topologies specify the connections and network information for the virtual machines and networking components within the emulation experiment. For this research, a total of seven topologies were created with varying security control mechanisms for the virtual DER network:

1. Flat network with plaintext Modbus traffic (no encryption)
2. Flat network with encrypted Modbus traffic using SSH tunnels
3. Segmented network with plaintext Modbus traffic (no encryption)
4. Segmented network with encrypted Modbus traffic using SSH tunnels
5. Critically Segmented network with plaintext Modbus traffic (no encryption)
6. Critically Segmented network with encrypted Modbus traffic using SSH tunnels
7. Flat network with plaintext Modbus traffic (no encryption) with a moving target defense overlay that randomized IP addresses every 3 seconds

These topologies use a combination of the various security measures mentioned in Section 2.2.3, with different enclaving strategies, encrypted tunnels, and MTD approaches between the

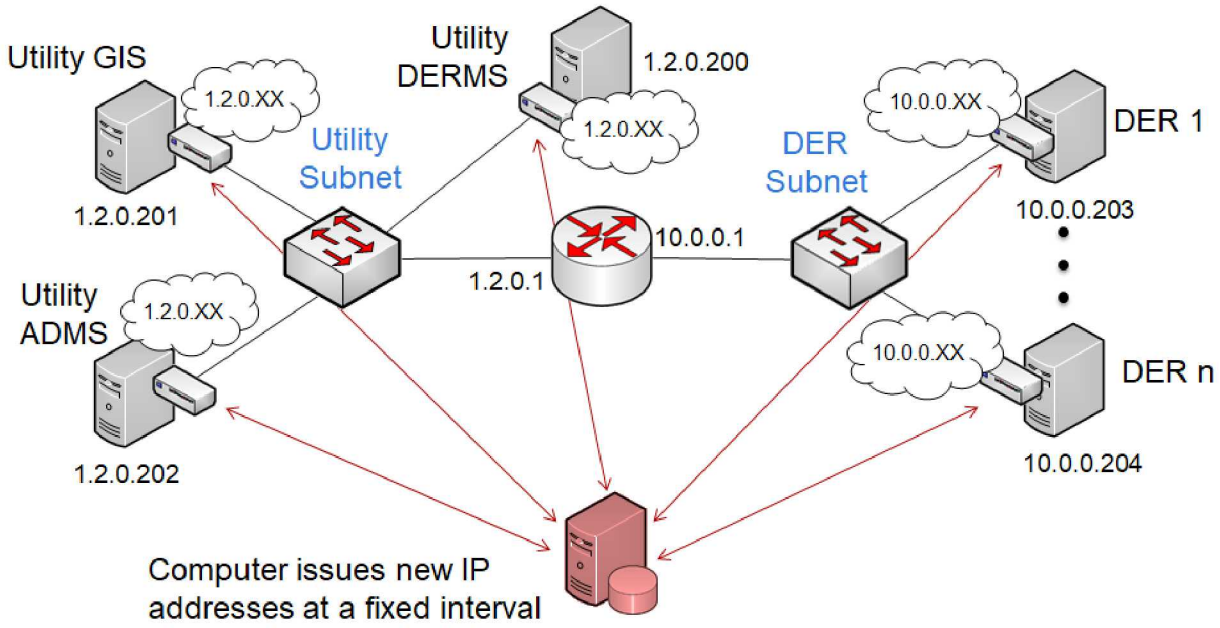

Figure 2.7. Implementation of Moving Target Defense on a DER communication network.

utility and DER subnets. In the case of topology #3, this environment was again deployed with simulated inverters and also replacing one of the inverters with a physical 3.0 kW residential-scale PV inverter with a HIL feedback loop.

Specifically, a VM running the SunSpec Validation Platform (SVP) [43] was placed in the utility network and was monitoring and controlling 20 DERs. The SVP was issuing voltage regulation setpoints according to the volt-var shift algorithm described in [31, 32].

Besides the various enclaving strategies, Each topology also simulated some extraneous traffic on the network through a tool in Minimega called Protonuke, which merely added some extra load by simulating internet traffic related to web browsing, mail, and SSH connections between extra VMs acting as servers and clients. Finally, a Kali virtual machine was deployed inside each topology for the red team to use in their assessment.

Flat Distribution Clear

The flat topology is a simple flat network for all twenty DERs in our system, meaning that each DER can reach and communicate with every other DER. Two routers were used to connect a utility network through a Wide Area Network (WAN), such as the internet or any other large distributed network. This configuration is shown in Figure 2.8. The flat distribution clear topology was intended to depict connection through a traditional ISP network, or similar, such as could be found in a large neighborhood or a public campus. The DER inverters were connected directly to the public network and shared that network with any other devices that would be connected to

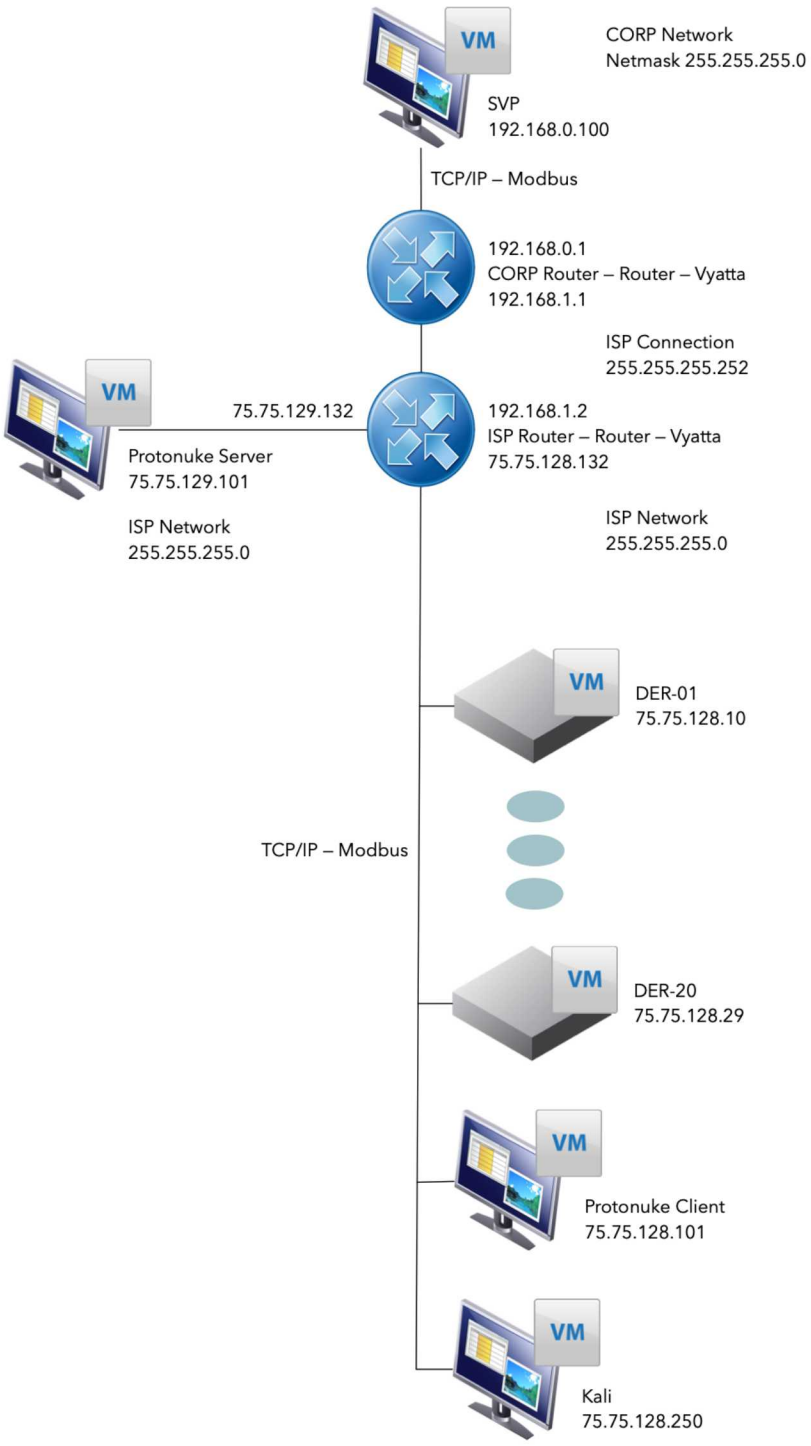

Figure 2.8. Flat Distribution Clear

it. In this topology, the SVP virtual machine connected directly to each DER inverter using TCP Modbus with no additional protections in place. Meanwhile, the Protonuke client shared the larger public network and connected to the Protonuke server on a separate network to generate some extraneous traffic as would be case on a used network. The Protonuke server was the sole device on this separate network and is only present to represent connections and traffic to and from other networks.

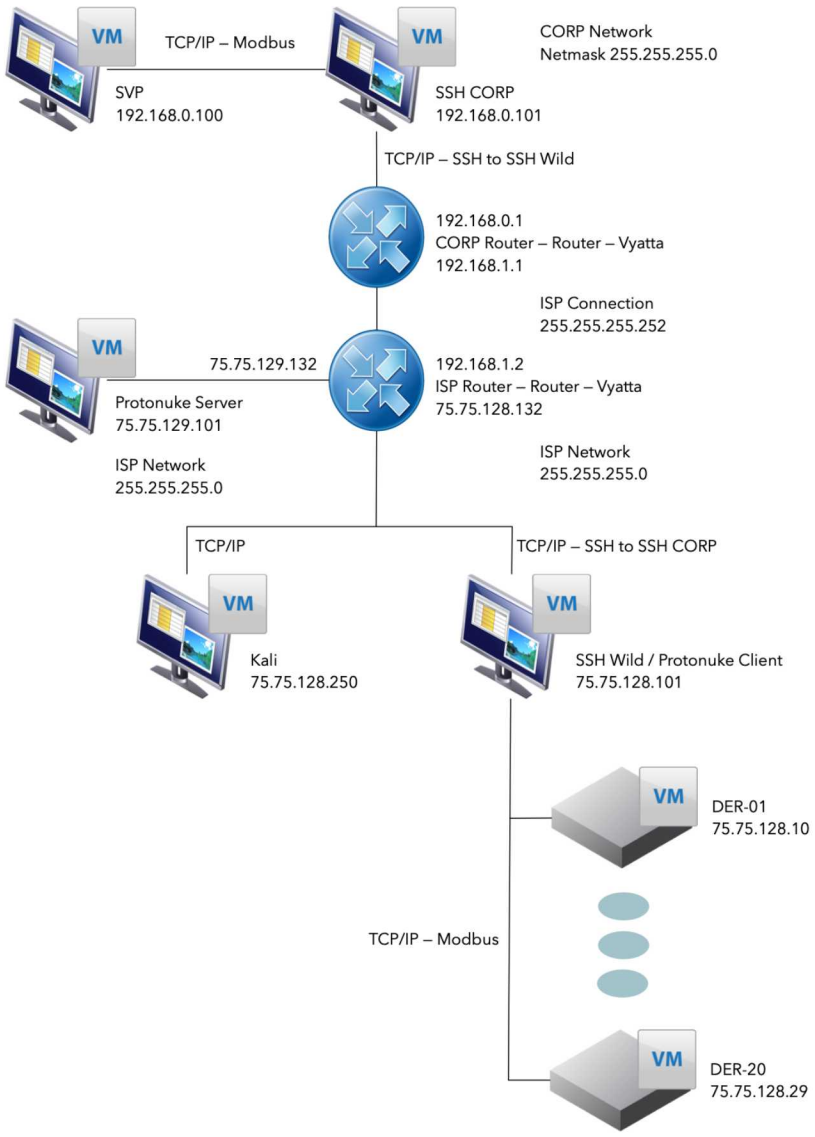

Figure 2.9. Flat Distribution Encrypted

Flat Distribution Encrypted

The flat distribution encrypted topology is very similar to the clear version described in the last section, but adds in an SSH tunnel as shown in Figure 2.9. This SSH forwarding tunnel was established to encapsulate the Modbus TCP traffic between the SVP virtual machine and each DER. This is performed by taking Modbus traffic in the clear on the unsecured sides of the "SSH CORP" and "SSH WILD" virtual machines and securing them using OpenSSH. In short, traffic within the flat control network for the DERs and within the "CORP network" are plain Modbus with no security, but across the public network these messages have been encrypted. This setup is meant to represent a utility providing an SSH server at the DER sites to be used to forward traffic through an encrypted tunnel. This limits exposure of the Modbus traffic but does not protect it all the way to each DER inverter. For instance, if an adversary could be positioned between the "SSH WILD" server and a DER inverter, they could capture and manipulate the Modbus traffic to the inverter with ease. An alternative to this topology would place a SSH server directly next to each DER inverter. The SVP system could then communicate over the public network to the SSH server at the DER, which would then communicate to the DER inverter directly. In that instance, an adversary would have to be on the private network to capture or manipulate the Modbus traffic, but this extra security comes with significantly increased costs due to the need for extra devices for SSH tunneling. This can be minimized by building security functionality, such as SSH tunneling, directly into smart inverters themselves.

Segmented Distribution Clear

In the segmented clear topology, the flat control network for the DERs are broken into three separate segments. The inverters themselves are split equally between these segments, with 6-7 inverters per segment as shown in Figure 2.12. By splitting the inverters into separate segments, the difficulty of impacting all twenty is increased. However, as this is a "clear" topology, each segment is directly routable and the Modbus traffic is passed in the clear and unsecured.

The segmented distribution clear topology was meant to break the DER inverters into separate, publicly routeable networks. This offers a couple things to the utility. First, it would remove the reliance on a single network to maintain traffic to all DER inverters. If there was a network outage on one segment, a portion of the DER inverters would still be available for communication with the ADMS. Additionally, this segmentation reduces the target surface for an adversary, who now requires access to each segment to intercept or manipulate Modbus traffic for that segment. In other terms, it increases the steps required to impact all the inverters. The Kali virtual machine and the Protonuke server were placed on a fourth network segment connected to the publicly routable network.

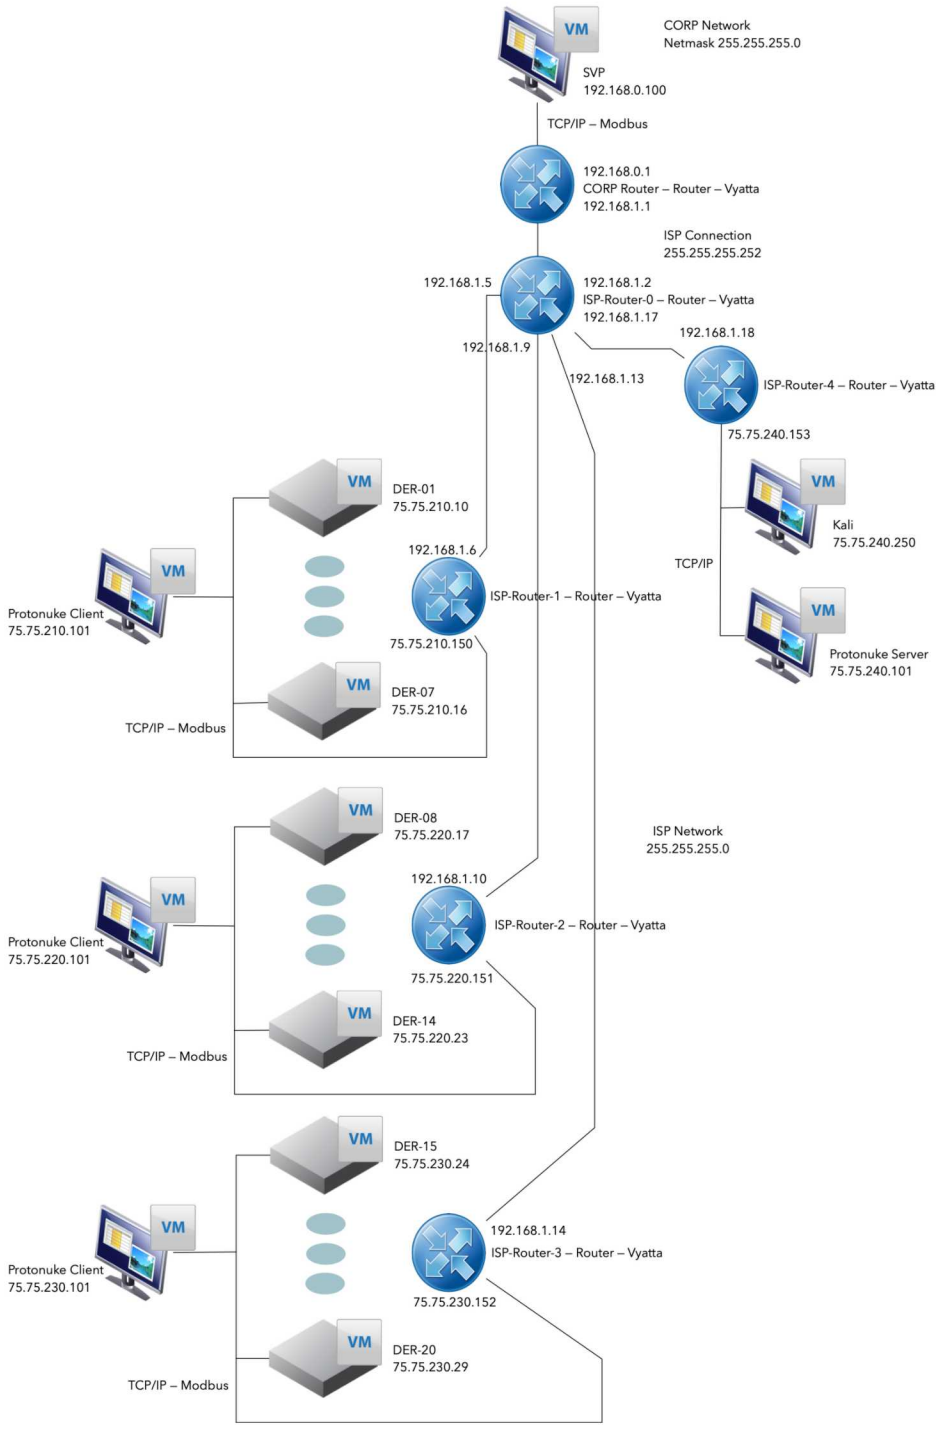

Figure 2.10. Segmented Distribution Clear

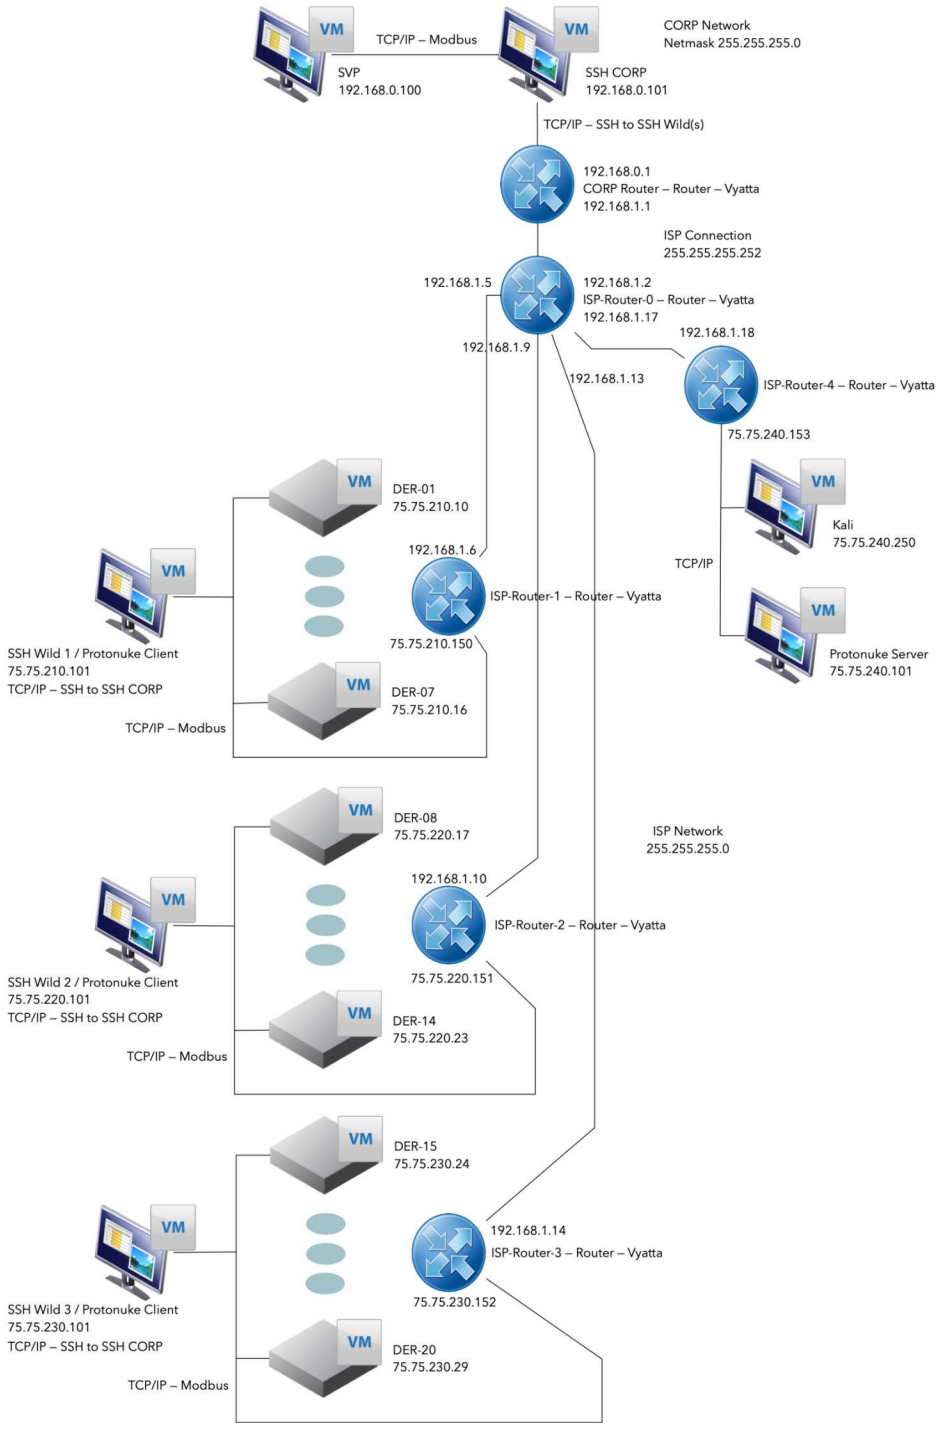

Figure 2.11. Segmented Distribution Encrypted

Segmented Distribution Encrypted

The segmented distribution encrypted topology was similar to the segmented distribution clear topology, with the main difference being the inclusion of SSH servers on each segment as shown in Figure 2.11. The SVP VM passed TCP Modbus traffic through the SSH CORP system and then through the encrypted tunnels with each of the "SSH Wild" servers (one per segment). Those SSH servers at each network segment then communicated in the clear to each DER inverter within their segment. As previoulsy noted for the flat encrypted topology, this will add protections to the Modbus traffic over the public network, but in this case, also comes with the additional measure of network segmentation to limit impact to grid support when some of the DERs are compromised.

Segmented Critical Distribution Clear

The segmented critical clear topology is very similar to the segmented clear topology, with the only difference being that the number of segments is increased to five, as shown in Figure 2.12. This means that each segment has only four inverters and thus the impact of a single segment being compromised is reduced. In other words, four inverters has less generation capacity than six or seven, so the impact of losing a segment is less than would be observed in the segmented topology. One other aspect that a segmented critical network topology brings into play is the ability to specify various assets as various levels of criticality. That is, some enclaves could be protected at higher levels than others depending on the corresponding system impact from loss of that enclave. For example, one segment could be prioritized for a hospital where another segment may not be as critical to maintain grid support capabilities. Again, since this is a clear topology the SVP control messages communicated with each DER inverter using TCP Modbus with no additional protections.

Segmented Critical Distribution Encrypted

Similar to the other encrypted topologies, the segmented critical distribution encrypted topology uses SSH tunnels to forward TCP Modbus to each inverter through an SSH server. Since there are now five network enclaves, there are five fielded SSH servers–one for each network segment as shown in Figure 2.13.

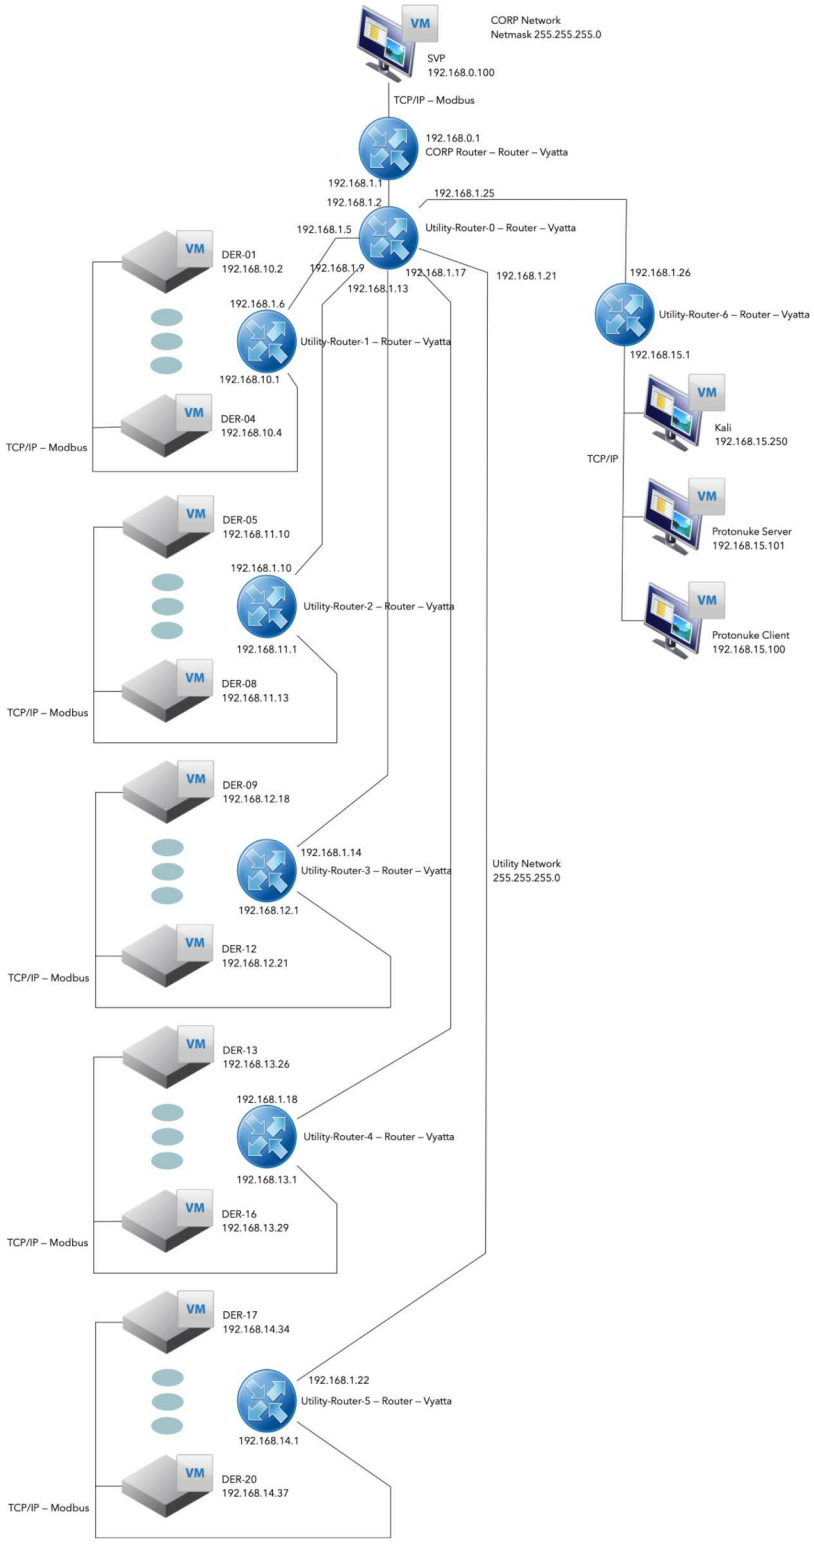

Figure 2.12. Segmented Critical Distribution Clear

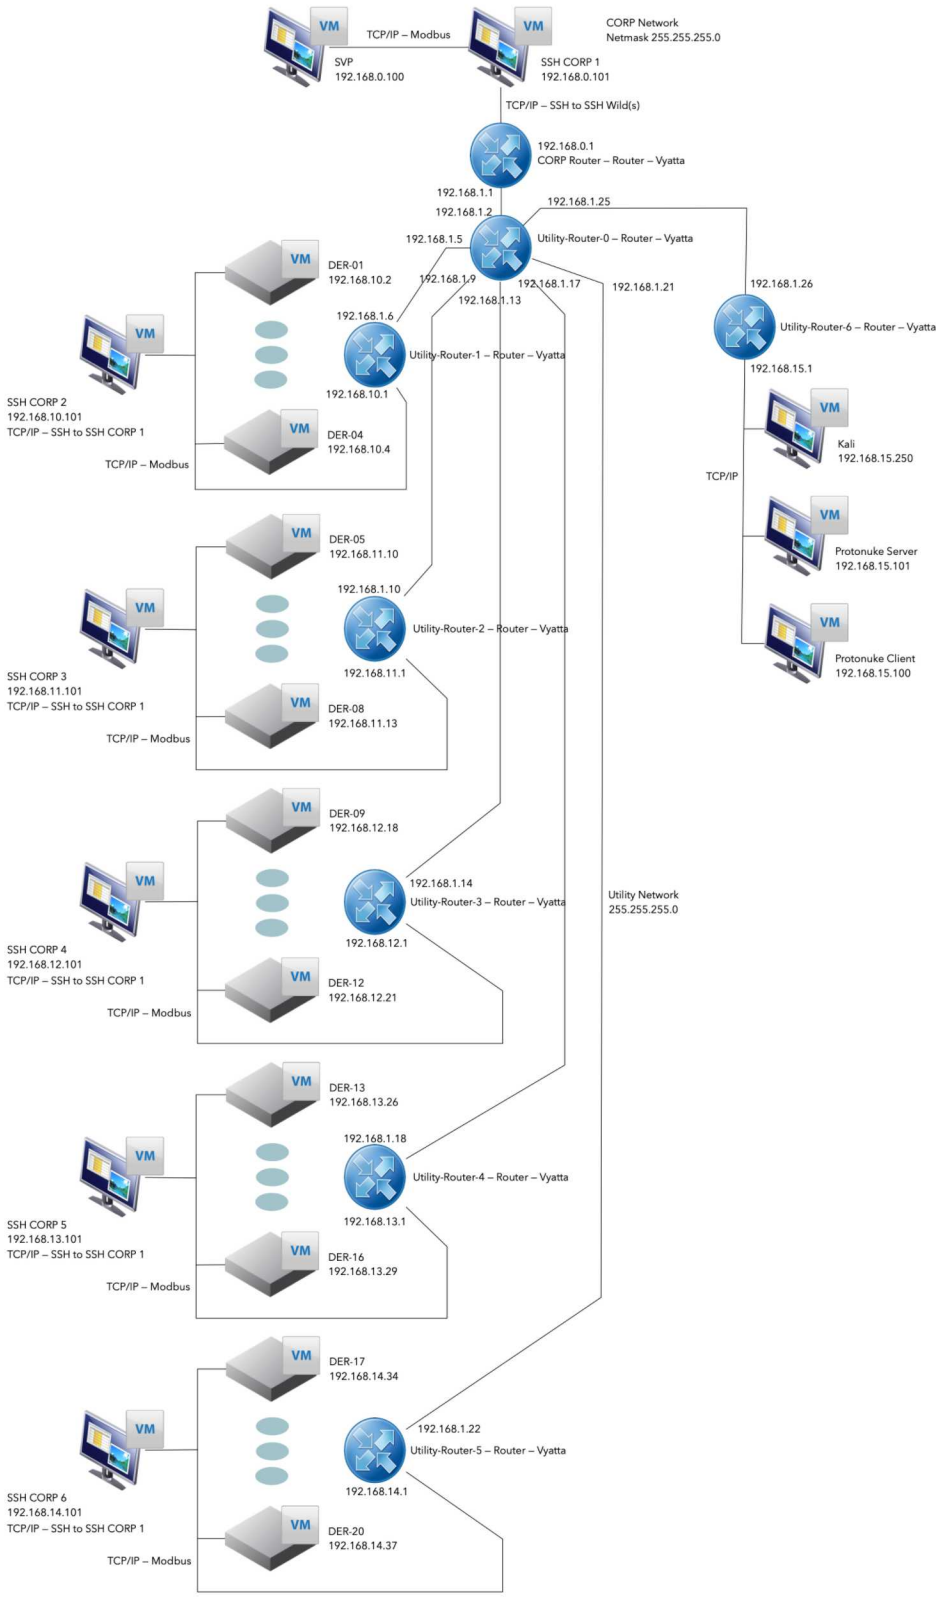

Figure 2.13. Segmented Critical Distribution Encrypted

This page intentionally left blank.

Chapter 3

Impact Assessment - Communication Latency

When cybersecurity features are added to control networks, there is an increase in communication latency from processing data, exchanging keys, binding certificates, performing encryption, or reconfiguring the system. These operations have the risk of adversely affecting real-time grid operations if the delays are significant. Several experiments were conducted to determine the communication latency associated with adding security features to DER networks. These studies were conducted using SCEPTRE and physical DER and Phasor Measurement Units (PMUs). Section 3.2 provides assumptions and caveats for the latency measurement experiments. Section 3.2 presents the impact to communication times as measured in several emulation experiments and Section 3.3 discusses the latency observed in several cases with real PMUs and smart inverters.

3.1 Latency Limits in Control Systems

Before continuing, it is important to note how the latency results obtained from a system emulation are useful and what the limitations are [44]. For instance, the absolute latency values for the various security methods studied will likely not be representative of hardware implemented in the field. Various implementations of the protocol stack and firmware and hardware variations may lead to very different results. However, relative impacts from applying additional security mechanisms are illustrative and help to provide understanding on the scale of the additional time required. This information is useful in determining whether system performance will or will not be significantly degraded. Field testing with components in the environment they will be utilized should still be conducted to verify operation is as desired.

While most of this chapter is presenting and examining results of communication latencies for various security mechanisms in a control network, this section quickly discusses limitations of communication latency experimentation. It is well known from the Nyquist–Shannon sampling theorem that when measuring transient behavior in a system, the sampling rate of the system measurements must occur at least twice as fast as the fastest behavior of interest. When sampling at rate greater than this limit, it is possible to reconstruct the behaviors of interest. The importance of this limit in this context is that it gives a bound for the sampling rate needed to observe power system dynamics. When communication latencies come into play, they should fall within those

bounds for the behaviors that are being observed or controlled, such as for analyzing transient stability. Likewise, in analyzing feedback control systems for DER, the sampling rate is important for the purpose of analyzing closed-loop stability of the system. Analysis of the closed-loop stability for various control mechanisms and behaviors of interest can include study of such as delay margin or the location of the discrete system poles (eigenvalues) demonstrate the regions of stability. Furthermore, power system concerns and constraints feed into the limits developed for the timing required for various controls in grid control.

As noted in DOE's 2017 report on the Modern Distribution Grid: Volume III, the communication timing requirements for DER are on the order of seconds, with typical bandwidth and latency requirements of 10 kbps and 5 seconds, respectively [45]. These communications requirements represent generalized limits on how much latency can be tolerated between the utility and smart inverters. Prior work on three transmission-level and one distribution-level distributed DER control algorithms provided a far more detailed view of the relationship between communication latency and performance DER control algorithms. It was found that the hierarchical volt-VAR shift distribution algorithm was effective with latencies up to 20 seconds [32], whereas the transmission services were severely impacted with lower latencies. Synthetic Inertia experienced a loss of machine synchronism defined by rotor angle separation with latencies between 200-400 ms (depending on the gain) [46]; communications-enabled fast acting imbalance reserve was ineffective if the delay is longer than the time to the frequency nadir (e.g., 1-10 seconds depending on system inertia) [47]; and communications-enabled DER droop control experienced oscillations with latencies of 110-400 ms (depending on the gain) [48]. These findings all indicate the control algorithm will lose effectiveness with increasing latency, leading to a range of potential problems.

3.2 Communication Latency (Emulated)

3.2.1 Network Segmentation

As expained above, the cybersecurity implications for power system performance depends on the grid-support service being provided. Certain communications-enabled DER services are robust to latency and other quality of service characteristics (dropped packets, DER availability, etc.), while others are not. In this section, a SCEPTRE experiment was created to calculate the increased latency associated with adding network segmentation. Notably, the main difference in these topologies is the extra hop from breaking the DER control network into multiple segments. A round trip time (RTT) for a segmented DER network and a flat network were calculated by pinging the DER from the utility (SVP) Windows VM. The results for more than 10,000 individual measurements is shown in Fig. 3.1. If we assume these results are normally distributed, the mean and variance of the distribution is shown in Table 3.1, along with the minimum, maximum, and median values. Overlayed on top of the histogram are plots of the normal distributions fit to the datasets. These distributions are scaled virtically to match the histogram and, thus, are not probability density functions (PDFs) because they are not continuous over time. Basically, the y-axis of this graph is a probability of a measurement falling into a bin of the histogram, not a probability of a measure-

ment being a certain value. Another important note is that these distribution results contain long tails on the right due to several outliers from instances when packets are dropped. This causes the probablity density functions (PDF) shown in Fig. 3.1 to appear as though they don't match the results well.

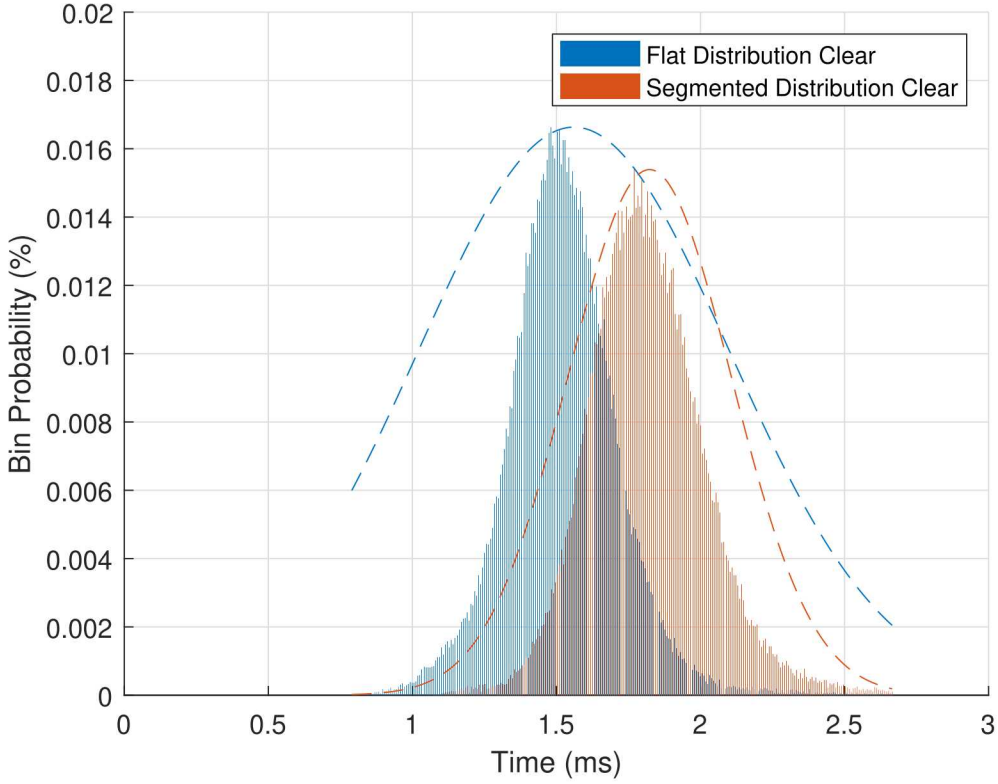

Figure 3.1. Differences in Communication Times for Flat and Segmented Networks using Modbus/TCP with no Transport Security

Table 3.1. Network Segmentation Latency using Modbus/TCP

Case	Mean, μ (ms)	Standard Deviation, σ (ms)	Min (ms)	Max (ms)	Median (ms)
Flat Distribution Clear	1.5605	0.5396	0.7861	16.8277	1.5192
Segmented Distribution Clear	1.8234	0.2834	1.0188	11.2763	1.8024

3.2.2 Encryption

Secure Shell (SSH) cryptography protocols were used to wrap unsecured Modbus communications at the transport layer. To show the impact of the encryption, a SCETPRE environment was

constructed with multiple SSH tunnels using common encryption cyphers and modes of operation. The network topology is shown in Figure 3.2. This experiment measured the communication time required for Modbus/TCP packets to traverse the network when wrapped in transport security using TLS. If we assume these results are normally distributed, the mean, standard deviation, min, max, and median of the distribution are as shown in Table 3.2. Note that due to the outliers (large maximum values), the standard deviation is far larger than one would expect from Figure 3.3. Note, there were a small number of cases where the packets did not reach the emulated inverters and the connection was reset verify this.

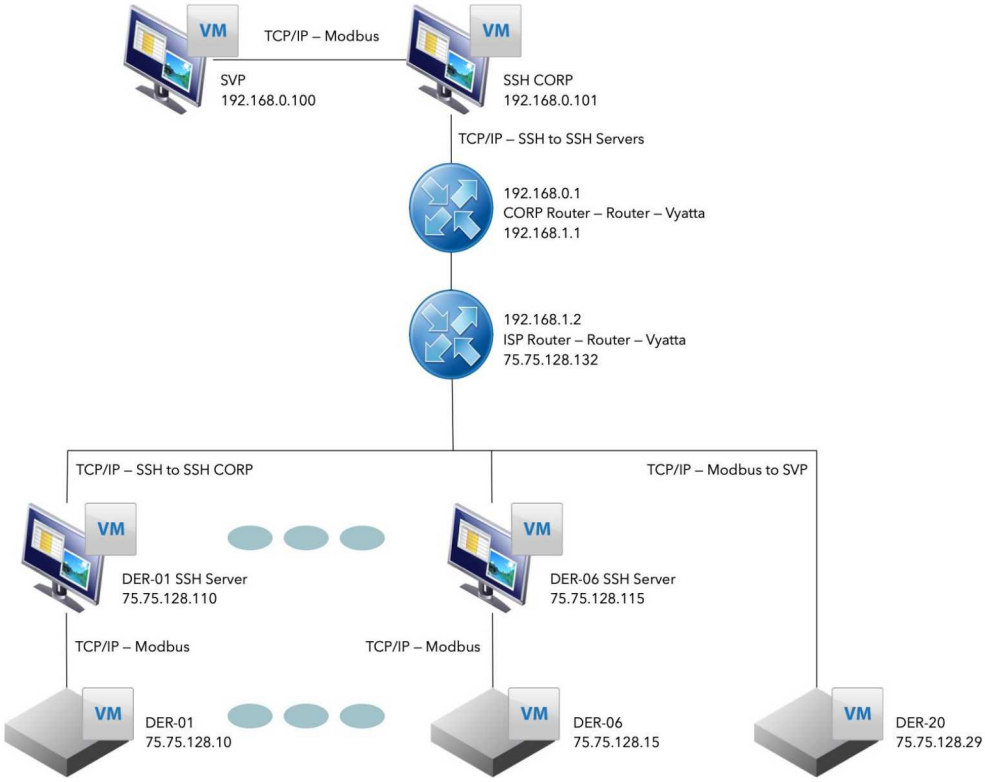

Figure 3.2. Topology for Testing Communication Latency of Encryption Ciphers and Cipher Modes with Transport Security

3.2.3 Moving Target Defense

Previous research on an emulated grid wide area network (WAN) has shown that ADDSec moving target defense can be beneficial to system security during a reconnaissance or denial of service type attack in which an attacker is sending packets over the network. In a study on ADDSec resilience [40], latency measurements in the form of round trip time were taken on a WAN in which one device has been compromised by a self-propagating worm. Without ADDSec, the network hosts were infected within minutes, leading to a doubling of latency. With ADDSec in operation, fewer

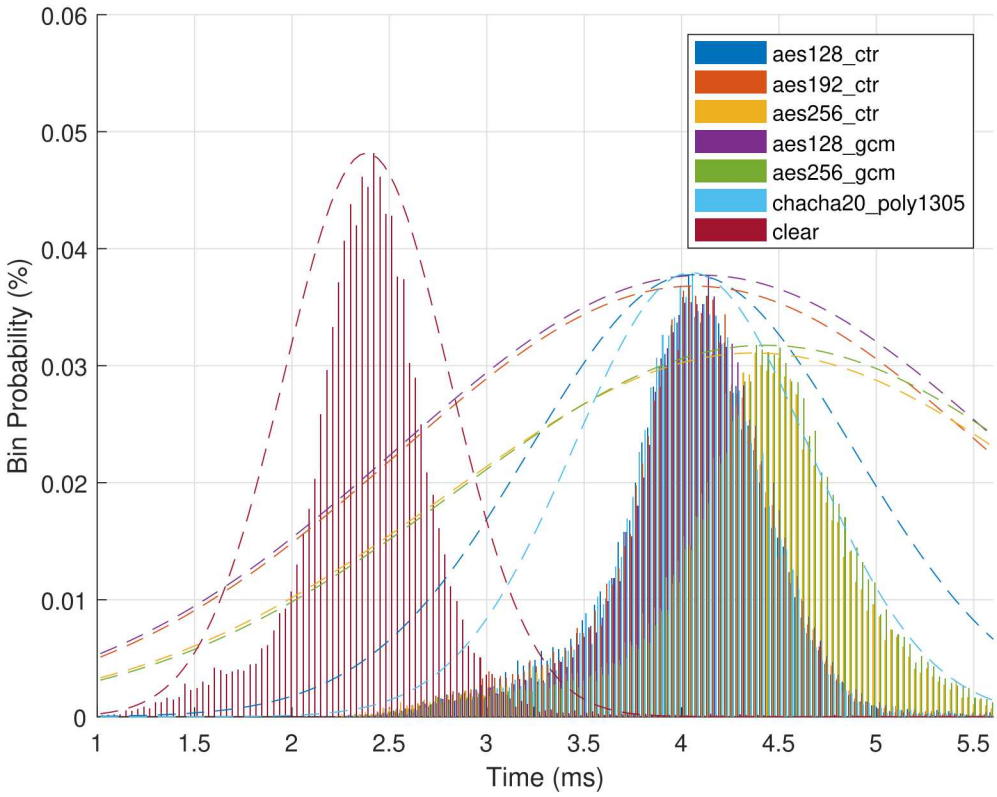

Figure 3.3. Differences in Communication Times for Common Ciphers and Cipher Modes for Transport Security of Modbus/TCP

Table 3.2. Encryption Latency using Modbus/TCP

Case	Mean, μ (ms)	Standard Deviation, σ (ms)	Min (ms)	Max (ms)	Median (ms)
AES128-CTR	4.0526	0.8295	2.0698	81.1382	4.0604
AES192-CTR	4.0662	1.5339	2.1778	206.7507	4.0748
AES256-CTR	4.3728	1.5879	2.0645	206.9327	4.3957
AES128-GCM	4.1056	1.5665	2.2905	205.8982	4.0985
AES256-GCM	4.4290	1.5858	2.2220	205.7683	4.4418
ChaCha20-Poly1305	4.0496	0.6043	2.1506	45.0614	4.0565
Clear	2.3834	0.4236	1.0010	15.1254	2.3847

network hosts were infected within the same timeframe, and latency was increased to a much lesser extent. Moreover, the study demonstrated that the overhead to network latency introduced by the ADDSec SDN controller during normal operations was minimal in comparison to the latency increase during an attack.

In applying MTD to a DER communication system, one must be careful to consider network constraints and the environment in which it is operating. Although modern DER grid-services do not have strict latency or timing requirements, this could change with the integration of more sophisticated transmission or distribution grid-support services. This said, the additional latency from MTD is nearly negligible. In prior work, communication latencies for various MTD modes were determined with different randomization time periods; it was found that MTD increased the average latency by less than 1 ms but caused slightly higher dropout rates (approx. 1 dropout per 33.3 seconds with IP randomization every 3 seconds) [40]. Other approaches to MTD, like path randomization, may increase latency more. An 11.73 ms increase in RTTs for path randomization was reported by Chavez [41]. However, even though MTD does not significantly increase latency on the system, it does potentially introduce other forms of system overhead that needs to be considered.

3.3 Communication Latency (Physical)

The following sections discuss results captured using physical hardware. First, the communication times for PMU messages between ABQ and several geographically distributed locations within the continental United States are discussed to better understand latency impacts of distance. Then, timing measurements for several smart DER are discussed based on tests conducted at the Distributed Energy Technologies Laboratory (DETL) at Sandia National Laboratories.

3.3.1 Geographic Separation

The results of Figure 3.4 show the transit times for one-way messages from the respective PMUs to Sandia National Laboratories in Albuquerque, NM. The PMU transit times to Albuquerque are calculated using the GPS timestamp and the GPS time at the receiver. The connection to Texas is over a dedicated fiber line and has minimal network hops. Conversely, the NM PMU has numerous routers and switches in the communication path which slow down the packets. In general, these results show that the architecture (switch and router hops) and communication medium (copper vs. fiber) is more important to data-in-flight times than geographic separation. This is important to keep in mind when developing the control network architecture to ensure that the number of network hops does not impair control system operations.

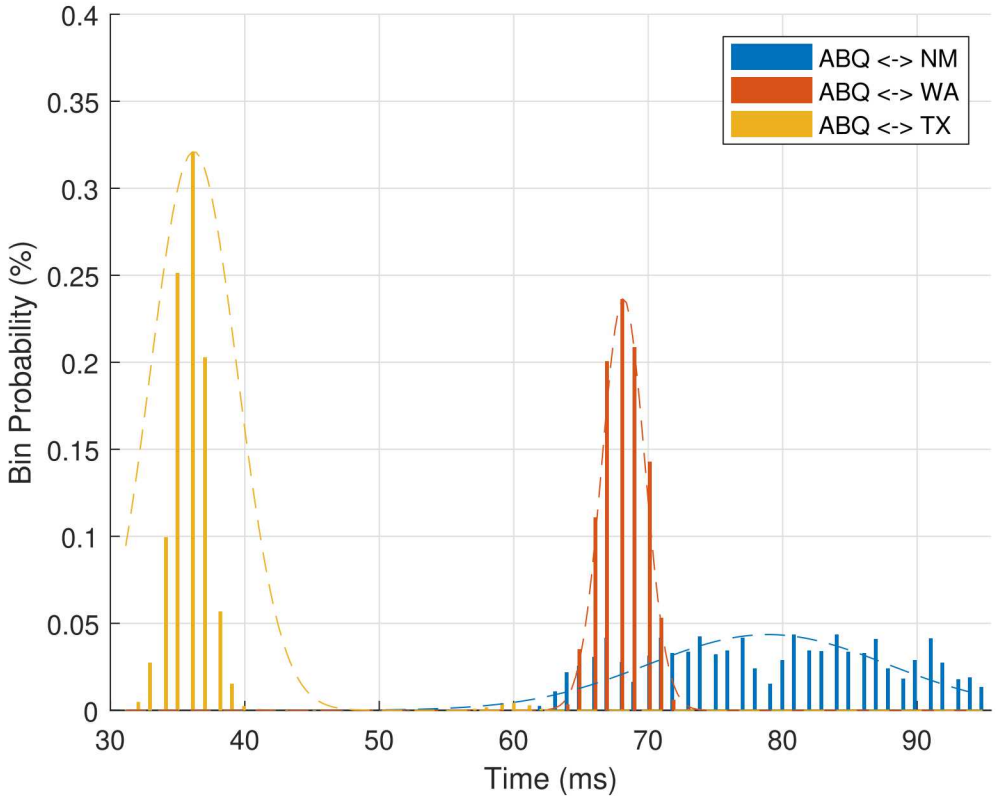

Figure 3.4. Differences in Communication Time from Geographically Separated Phasor Measurement Units to ABQ

Table 3.3. Encryption Latency using Modbus/TCP

Case	Mean, μ (ms)	Standard Deviation, σ (ms)	Min (ms)	Max (ms)	Median (ms)
PMU-1 (ABQ $\leftrightarrow$ NM)	78.9117	8.9063	61.0000	105.0000	79.0000
PMU-2 (ABQ $\leftrightarrow$ WA)	67.1551	1.5846	63.9999	86.0002	68.0001
PMU-3 (ABQ $\leftrightarrow$ TX)	36.2080	3.2368	30.9999	66.9999	36.0000

3.3.2 Smart Inverter Read and Write Times

1000 Modbus read and write times were collected for two commercially available residential-scale DER devices and one CHIL device [49] in DETL using the SunSpec System Validation Platform (SVP). The results are shown in Fig. 3.5 and Table 3.4. Inverter 1 has a large standard deviation for both read and write times. It is not clear if there are internal communication checks or other inverter processes that could be slowing some of the responses. This connection also includes a large number of outliers which significantly affect the distribution. Similar results are reported in [38]. Like Inverter 1, Inverter 2 had a direct connection of Modbus/TCP over 1 network hop, but responds much faster to both read and write requests. The connection to Inverter 3 included an Ethernet-to-Serial converter in the path to translate Modbus/TCP to serial Modbus. This added an additional delay due to the processing required to perform that conversion–possibly accounting for some of the larger average communication times for reads and writes with that device. It is believed that the variations observed in these results are not primarily due to the various network architectures. It is more likely that the inverters include different implementations of the protocol stack, processor hardware, and scheduling differences of processing tasks for I/O to and from memory. Further analysis would be required to determine the precise reason for the variations and draw generalization about the expected DER read and write times.

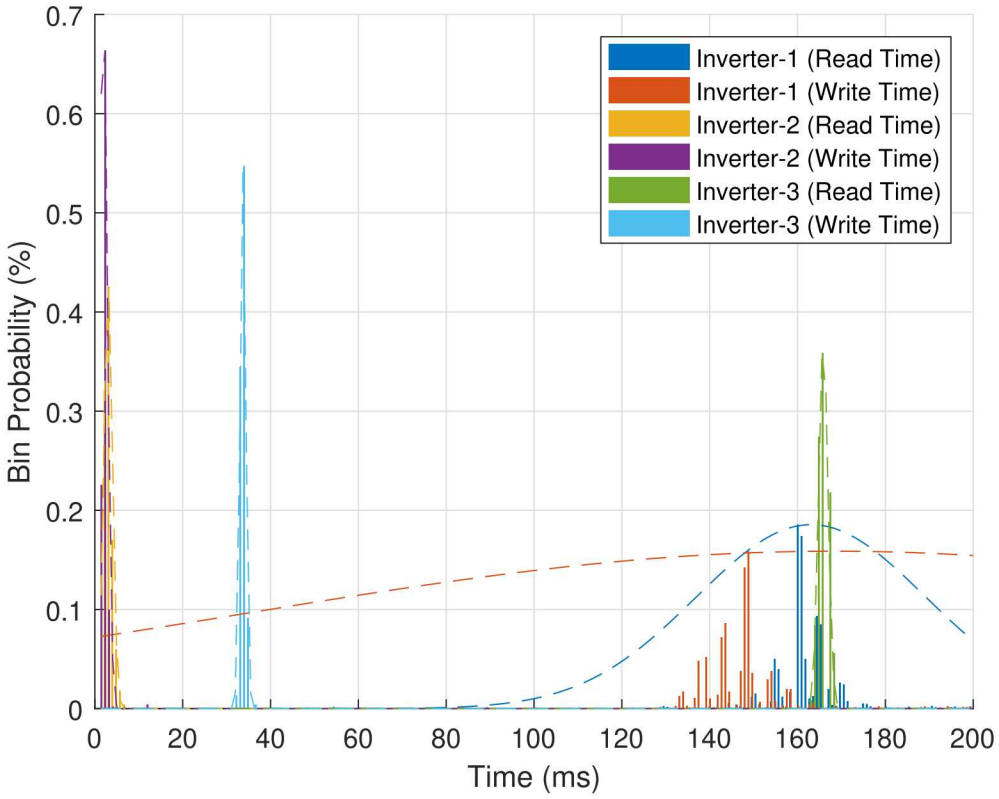

Figure 3.5. Differences in Communication Times for Several Common Smart Inverters

Table 3.4. Round-trip Communication Time for Modbus/TCP with Smart Inverters in DETL

Case	Mean, μ (ms)	Standard Deviation, σ (ms)	Min (ms)	Max (ms)	Median (ms)
Inverter-1 (Read Time)	163.0757	26.1437	44.9998	1145.9999	161.0000
Inverter-2 (Read Time)	3.0319	0.9801	0.9999	7.0000	3.0000
Inverter-3 (Read Time)	165.8618	1.0560	162.9999	168.0002	165.9999
Inverter-1 (Write Time)	168.3799	133.6979	38.0001	1435.0002	148.0000
Inverter-2 (Write Time)	1.9383	0.9110	0.9999	12.0001	2.0001
Inverter-3 (Write Time)	33.7298	0.6583	31.9998	36.0000	33.9999

3.4 Latency Observations

Based on the results for network segmentation, encryption, MTD, geographical separation, and DER read/write times, some observations can be made about the impact to the control system when adding security features. In general, large geographic distances have the possibility of adding 50-100 ms of latency for utility-to-DER communications due to the additional networking equipment (routers and switches) between endpoints. DER read and write times vary widely; they can be 1 second or larger in some situations. In contrast, network segmentation adds less than 1 ms, encryption adds on the order of 3-5 ms of additional latency, and MTD adds 1 ms. Therefore, for the proposed cybersecurity features, it is not believed they will impact the grid-support service performance since they adds only contribute a minor percentage of the total latency between the utility and DER.

This page intentionally left blank.

Chapter 4

Security Assessment - Red Teaming

Chapter 3 discussed the impact of DER network latency on various grid-support controls. This chapter will quantify the impact of networking security features on the security of the system. This is done through red teaming experiments with the goal of assessing the performance of each network topology in Chapter 2 under a range of attack categories.

4.1 Red Teaming Approach

Red team assessments are authorized, adversary-based, cyber assessment conducted to strengthen defenses through awareness and exploitation of the system's potential vulnerabilities. The primary objectives for each of the topologies assessed was to identify and compromise the DER devices (power inverters) by modifying communication or grid-support functions (Freq-Watt, Volt-Var, Power-Factor settings, etc.) or disrupting network communications.

4.1.1 Scope and Rules of Engagement

The security assessment focused on the communications between the emulated DER devices on the network, the simulated corporate and provider networks and any hardware in the loop (HIL) devices. Rules of Engagement were defined as:

- Limited to the SCEPTRE experiment network
- HIL devices are in-scope
- SCEPTRE, Phenix, Minimega, OpenDSS, PowerWorld are out-of-scope

4.1.2 Methodology

The assessment incorporated elements from Sandia National Laboratories' Information Design Assurance Red Team (IDART), NIST's Guide to Industrial Control Systems (ICS) Security Guidelines, Department of Homeland Security's ICS-CERT Recommended Best Practices, and collective

expertise of PV inverter systems. The red team assessed each of the topologies with a specialized methodology based on this guidance.

The assessment team developed impact metrics based around the triad of confidentiality, integrity, and availability, also known as CIA, which serve as the core attributes for many cybersecurity risk evaluation frameworks, including the NIST Common Vulnerability Scoring System (CVSS) Impact ratings [50]. These attributes are prioritized according to the system environment and mission, with relative importance levels captured in a system critical matrix (SCM). In this set of experiments, the DER network mission is constant, while the environmental parameters change in accordance with the network topologies.

4.1.3 Tools

Reconnaissance of each network topology began by actively probing from the Kali machine and other Linux machines on the subnetworks provided. The use of Nmap and OpenVAS provided IP identification, host fingerprinting and vulnerability assessments of the devices on the network. Tcpdump and Wireshark were utilized on the networks to capture packets on the wire for use in replay as well as for identification of communication protocols for modification and fabrication attacks. The use of open-source tools, SunSpec Dashboard, vendor-specific applications, and Simply Modbus were used to craft vendor specific protocol traffic to the devices on the networks.

4.1.4 Emulytics Challenges

Emulytics^{TM} environments enable rapid security prototyping and red teaming exercises. The networks were designed to represent realistic network topologies and passed real DER protocol and encryption packets. To further improve the fidelity of the derived results, a power hardware-in-the-loop (HIL) PV inverter was added to enable better system tests augmented with actual physical systems. However, while Emulytics^{TM} environments do well in faithful simulations of cyber-physical systems, the attack surfaces are typically reduced; human elements are removed, hardware, software, and firmware diversity are decreased, and overall emulated system complexity is limited. In some instances, the discovered vulnerabilities may be artifacts of the testbed setup itself, as they may be introduced by the emulation and not present in the field. The biggest challenge was found to be the interactions between the backend processes - SCEPTRE, Phenix, Minimega, OpenDSS, PowerWorld becasue they were a disparate set of tools that had not been designed to interface together. Due to these limitations, emphasis is placed on the rules of engagement which define the scope of assessment and were designed to have the red team concentrate on realistic vulnerabilities.

Although the red team methodology defined criticality levels and quantitative scoring values, they are still largely subject to the prior experience and priorities of the assessor. Assessors with varying levels of familiarity with a particular type of DER device, protocol, or system architecture may assign varying criticality levels to the same information compromise based on their interpretation of the potential impact, and subject matter expertise in DER operations is needed to accurately

grade consequences of compromise to the system itself. Moreover, resources are measured in the form of time to compromise, which is subject to variability with human actors in the loop and variation in the system setup. This adds variability and uncertainty in the results that are difficult to remove. Even autonomous red teaming systems, which are meant to produce reproducible baselines, are currently still reliant on feedback from human expertise. Given these drawbacks, the red team's quantitative risk analysis would be informed based on the frequency and ease of common attacks and the goals of this assessment.

4.1.5 Threat Catalog

The red team developed a threat catalog of vulnerability tests based on the goals of the asssesment, to categorize the types of vulnerabilities that a threat actor may seek to exploit. These Internet Security tests are listed below.

- Network Surveying
- Port Scanning
- System Identification
- Services Identification
- Vulnerability Research
- Router Testing
- Firewall Testing
- Password Cracking
- Denial of Service Testing
- Network Surveying
- Port Scanning
- System Identification

The red team modeled a threat from an attacker equipped with specialized knowledge of DER system protocols with considerations for insider access. The following threats were examined and executed during the assessment:

- Data Compromise: alteration or access of confidential of data by unauthorized users.
- Remote Exploits: exploiting existing privileges for authorized users on the system.

- Local Exploits: exploiting known CVEs in user applications.
- Interception: man-in-the-middle (MITM) or eavesdropping of authenticated communciations.
- Denial of Service: rendering the system unusable to authorized users, such as overloading the RTU processors.
- Policy: exploiting flaws in policy, such as firewall security settings.
- Insider Threat: exploiting authorized user knowledge or access for malicious purporses.

By tailoring each assessment against this catalog, the red team ensured their methodology was reproducible and applicable across a wide range of systems and threat models.

4.2 Results

To execute the assessment, the red team used network reconnaissance and network attack tools on Linux and Windows OS. The red team conducted assessments for two scenarios:

- **Outsider** (Public Network Attacker) An intruder who does not have access to a local subnet where the inverters are deployed. This adversary has no access to the DER device but does have access to one of the ISP routers.
- **Insider** (Local Attacker) The intruder is on the DER home area network (HAN) with a foothold on the subnet.

For each topology, the team conducted reconnaissance and active attacks including Packet Replay, Denial of Service (DoS), and Man-in-the-Middle (MITM).

4.2.1 Reconnaissance

Network scans using Nmap and OpenVAS discovered and fingerprinted devices. Nmap was used to discover devices and networks that existed in the topology. Figure 4.1 shows the results from an nmap scan. Nmap was run at different levels of granularity to discover open ports and determine basic OS fingerprinting. Figure 4.2 shows the open ports on the CORP Network (Windows SVP) machine and the Protnuke server (ISP network).

OpenVAS was then used to probe the open ports and test for vulnerabilities. Figures 4.3 and 4.4 show the scan result from an inverter and details on a vulnerability, respectively.

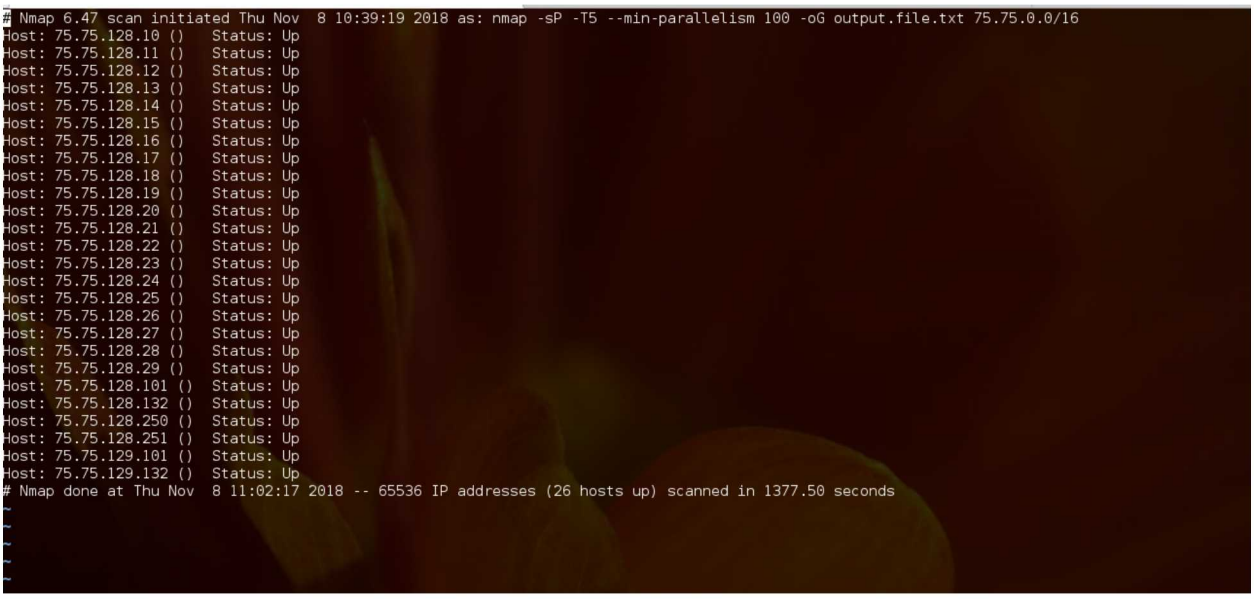

Figure 4.1. Nmap host discovery scan

SunSpec Dashboard application is designed to communicate with Modbus SunSpec RTUs. Access parameters required to connect are IP address, IP port, slave ID, and timeout period. The application displays available registers on the device and, depending on the parameter, may be writeable. Figure 4.5 shows the connection to a SCEPTRE RTU inverter.

By default, Modbus slaves listen on port 502. From the results of the above scans, Modbus was identified to be running on a custom port. The inverters were easily accessible with the SunSpec Dashboard application on this port and inspecting the connection revealed the SunSpec ID number, 0x53756e53, identifying the SunSpec Modbus Map for the Modbus traffic on the non-standard port 5502, shown in Figure 4.6.

Wireshark was used to analyze and reverse-engineer the communications using SunSpec Dashboard. For example, the mappings of SunSpec Model 101 (0x65), length 50 (0x32) are shown in Wireshark in Figure 4.7.

The settings for each Modbus register address were mapped by capturing packets for each value read, as seen in the packet capture. For example, the Modbus Nameplate registers were read using SunSpec Dashboard and packets were captured similar to Figure 4.7 above, and Inverter Model 101 bytes are shown in Figure 4.8.

4.2.2 Packet Replay

Packet replay is an attack in which data transmission is resent or repeated in a manner that causes undesired results. Utilizing the mapping discovered using Wireshark and the SunSpec Dashboard,

```
root@       :~# nmap -O 192.168.0.100

Starting Nmap 6.47 ( http://nmap.org ) at 2018-09-10 08:03 MDT
Nmap scan report for 192.168.0.100
Host is up (0.0024s latency).
Not shown: 991 closed ports
PORT      STATE SERVICE
135/tcp   open  msrpc
139/tcp   open  netbios-ssn
445/tcp   open  microsoft-ds
49152/tcp open  unknown
49153/tcp open  unknown
49154/tcp open  unknown
49155/tcp open  unknown
49156/tcp open  unknown
49157/tcp open  unknown
Device type: general purpose
Running: Microsoft Windows 7|2008
OS CPE: cpe:/o:microsoft:windows_7::- cpe:/o:microsoft:windows_7::sp1 cpe:/o:microsoft:windows_server_2008::sp1 cpe:/o:microsoft:windows_8
OS details: Microsoft Windows 7 SP0 - SP1, Windows Server 2008 SP1, or Windows 8

OS detection performed. Please report any incorrect results at http://nmap.org/submit/ .
Nmap done: 1 IP address (1 host up) scanned in 15.38 seconds


root@       :~# nmap -O 75.75.129.101

Starting Nmap 6.47 ( http://nmap.org ) at 2018-09-10 08:10 MDT
Nmap scan report for 75.75.129.101
Host is up (0.00079s latency).
Not shown: 996 closed ports
PORT    STATE SERVICE
22/tcp  open  ssh
25/tcp  open  smtp
80/tcp  open  http
443/tcp open  https
Device type: general purpose
Running: Linux 3.X
OS CPE: cpe:/o:linux:linux_kernel:3
OS details: Linux 3.11 - 3.14
Network Distance: 2 hops

OS detection performed. Please report any incorrect results at http://nmap.org/submit/ .
Nmap done: 1 IP address (1 host up) scanned in 15.21 seconds
```

Figure 4.2. Nmap host fingerprinting results

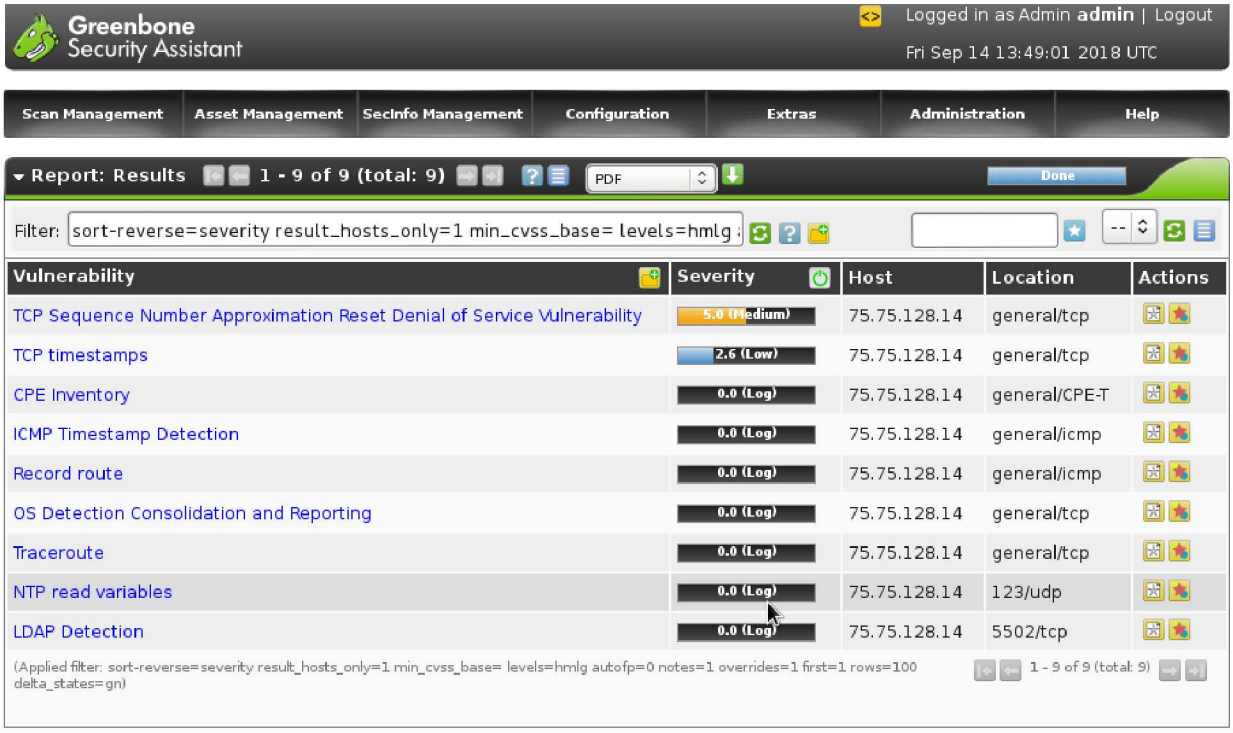

Figure 4.3. OpenVAS scan result on an inverter

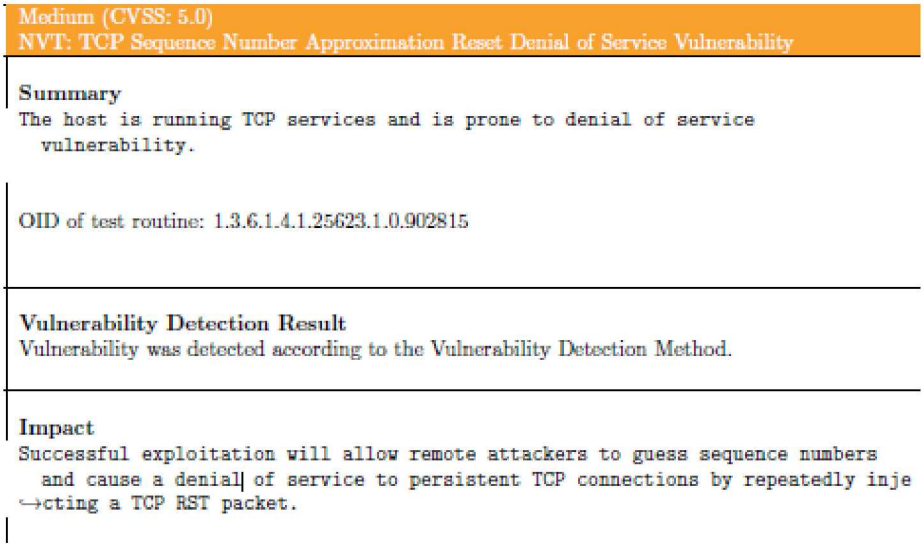

Figure 4.4. OpenVAS vulnerability description

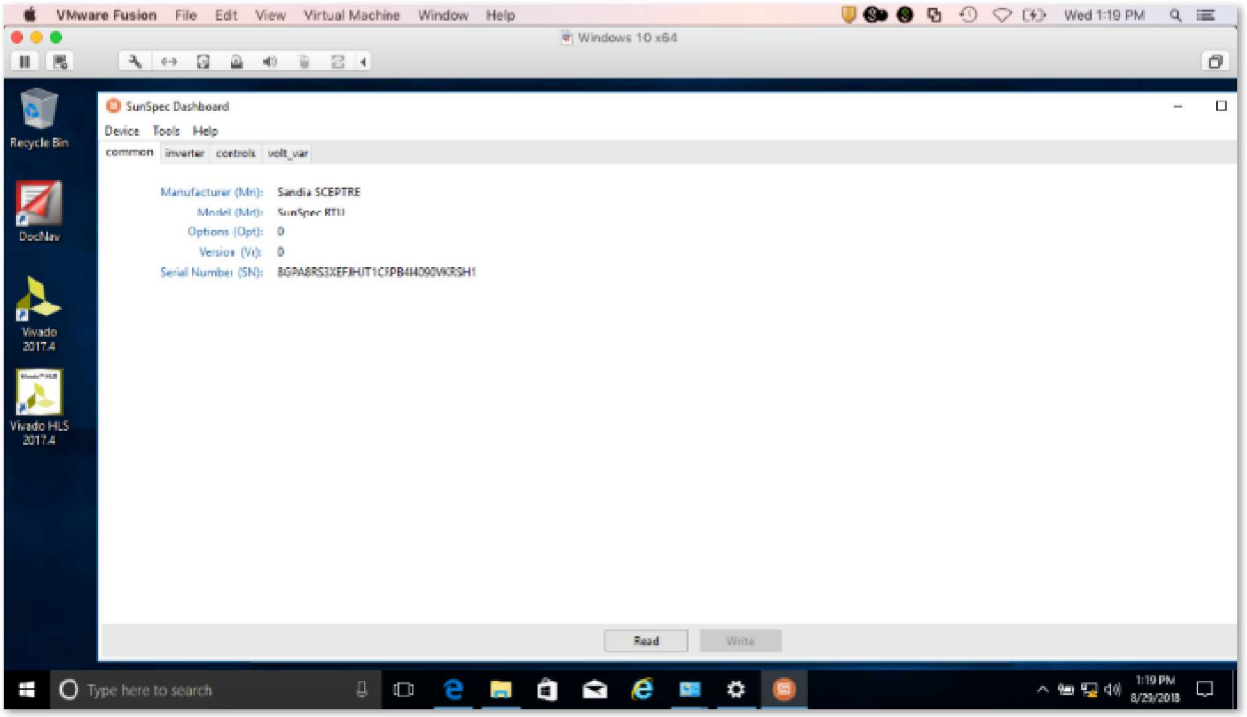

Figure 4.5. SunSpec Dashboard connected to an inverter

packets captured in the reconnaissance could be resent via Netcat, a Linux network communication tool. With some value modification and scripting, packet replay attack could be converted to a fabrication carried out on all inverters and multiple register locations autonomously. Shown in Figure 4.9 below is the SunSpec Dashboard showin inveter connection being cycled off via a script.

Other replayed packets included modification of inverter phase voltages, DC voltages, current, and power. Unauthorized actions demonstrate that an adversary could easily transmit fraudulent data to falsify the inverter's state and disrupt network communications.

4.2.3 Denial of Service

A Denial of Service (DoS) attack is a network attack in which data transmissions are used to render a system unavailable to legitimate users. Reconnaissance of the inverters indicated that they were susceptible to TCP SYN Flood attacks. Tools in the Kali suite called floodrouter6 and hping3 were used to send spurious router advertisements and TCP SYN requests from random IP addresses to devices on the network, respectively, causing network communication outages. A successful DoS attack preventing a legitimate user from accessing an inverter is shown in Figure 4.10 below. The DoS attacks were successful on the inverters and network devices, and between inverters and the SVP, in most cases.

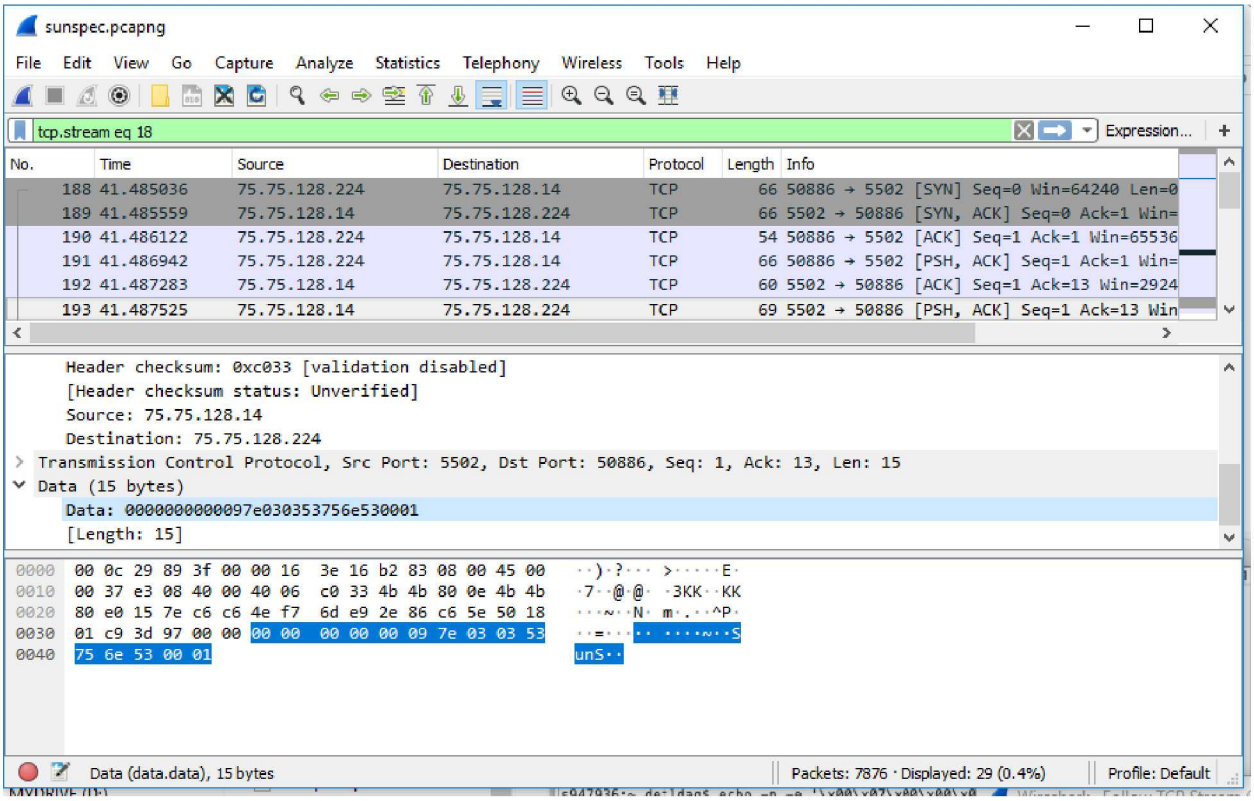

Figure 4.6. SunSpec ID number on port 5502 using SunSpec Dashboard

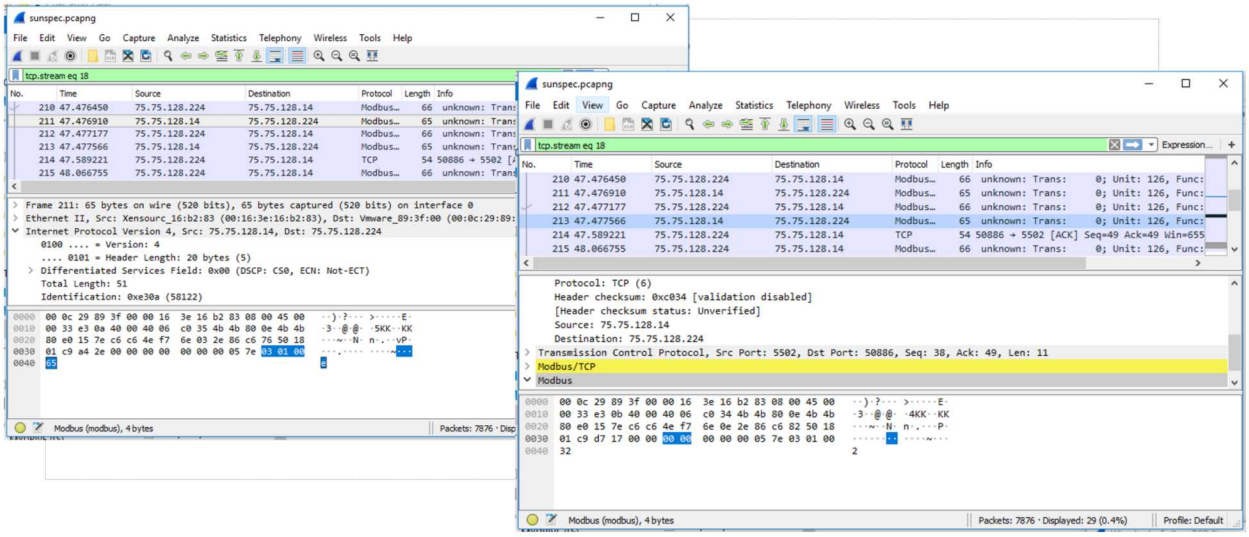

Figure 4.7. SunSpec Table 101 (0x65), length 50 (0x32)

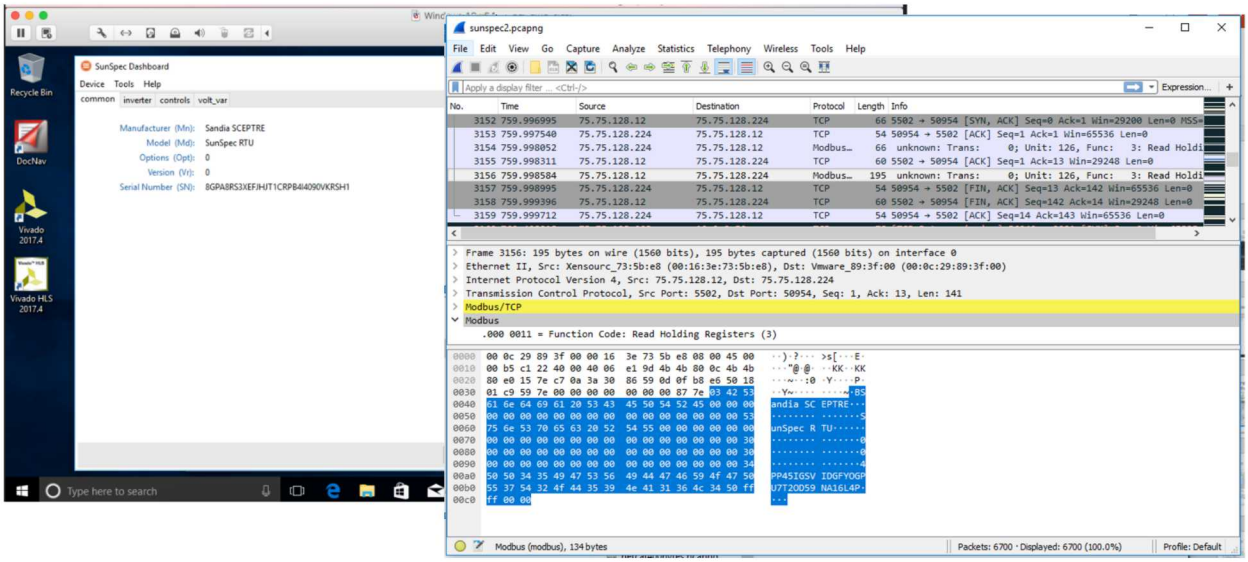

Figure 4.8. Mapping of Modbus Nameplate registers

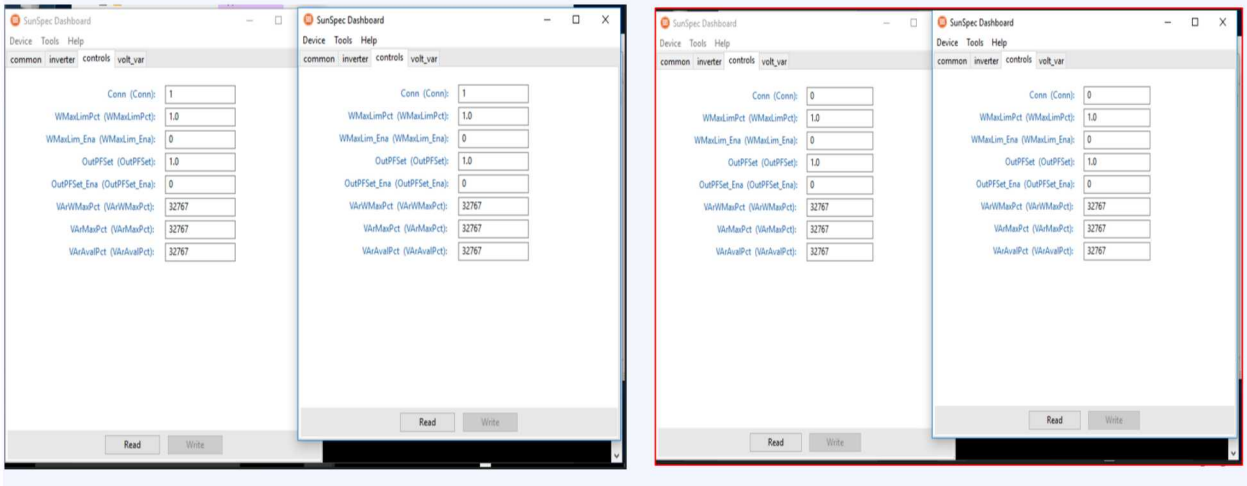

Figure 4.9. SunSpec Dashboard showing the connection register state cycling on two inverters

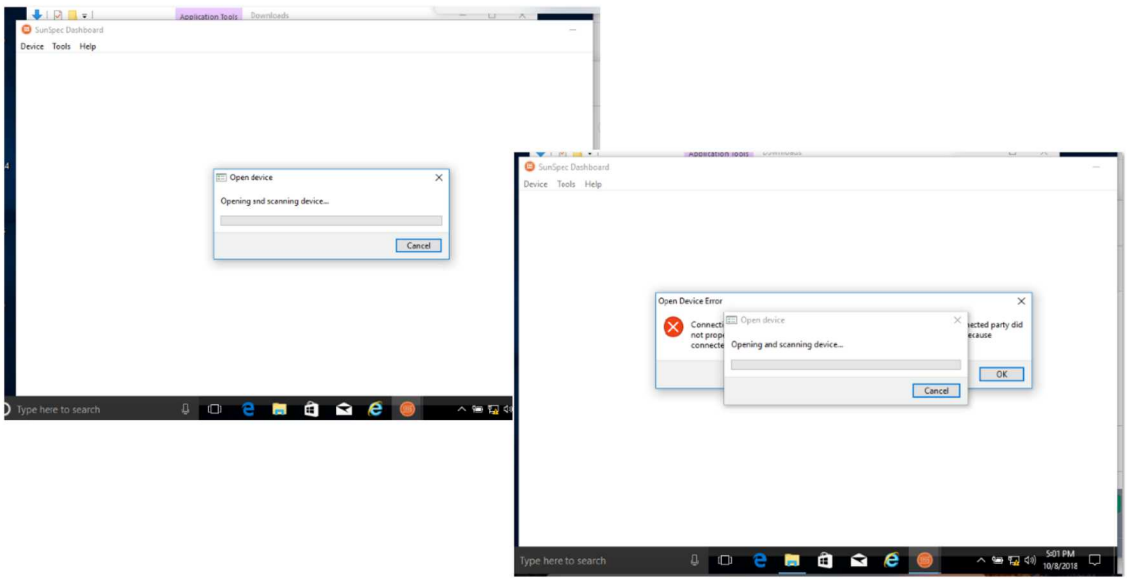

Figure 4.10. Initial inverter connection and unsuccessful inverter connection during a DoS attack

4.2.4 Man-in-the-Middle

Man-in-the-Middle (MITM) is the act of eavesdropping, dropping, delaying or altering communications while in transit from source to destination. For this attack, a tool in the Kali suite called Ettercap was used. The tool ARP cache poisons two devices of which the communications are desired to be intercepted. In all cases, MITM attacks worked for devices on the same subnet. In unencrypted topologies, MITM attacks between an inverter and the SVP saw the attacker stand between the inverter and its gateway router to capture Modbus packets which were visible in plaintext. Figure **??** shows eavesdropped Modbus/TCP traffic on the network.

4.2.5 Flat Network Topology without Encryption

- **Observations** Reconnaissance showed the Red Team's position was on the same subnet as the inverters. A router separated the utility's DERMS system into a separate network. The router and DERMS were susceptible to Denial of Service attacks. Man-in-the-Middle attacks were possible between each inverter and the router. Packet replay was possible directly to the inverters.

- **Challenges** As a baseline for the assessment, no challenges were found.

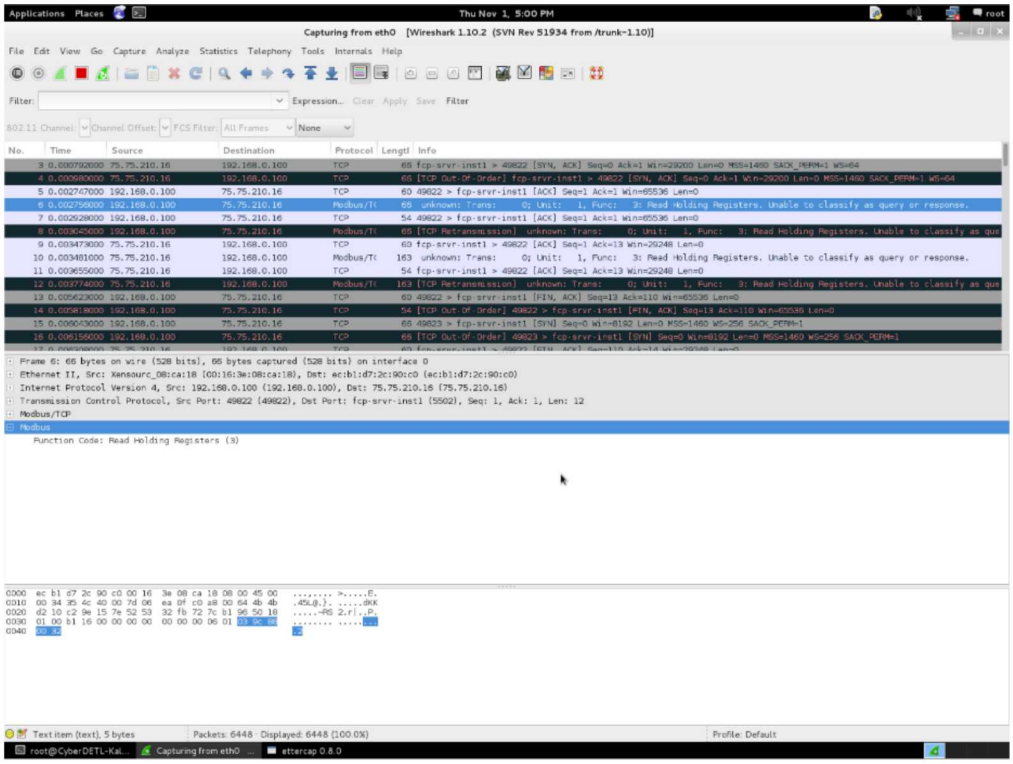

Figure 4.11. Sniffed Modbus/TCP traffic on one of the subnetworks

4.2.6 Flat Network Topology with Encryption

- **Observations** Reconnaissance showed that encryption was added via a bump-in-the-wire (SSH server) technique. However, the traffic on the subnet was unencrypted on the subnet where the inverters resided until it passed through the SSH server. The implemented architecture differed from the intended design shown in Fig 2.9, in that the DERs were not successfully deployed behind the SSH server. The DERs were instead immediately connected to the ISP Router, just as the Kali and SSH VMs. On this flat nework, the Red Team's tools were able to reach directly to the inverters, which allowed for register changes using Netcat and SunSpec Dashboard without challenge. In this, and all encryption enabled topologies in this assessment, the SSH gateway machines had a password-less root login enabled for SSH, an oversight from the development team setup. The Red Team was able to log in, pull SSH encryption keys and fingerprints, capture traffic from the inverters before encryption, and pivot onto the corporate segment of the network through the SSH tunnel. On this topology, DoS, MITM, and packet replay were all successful.

- **Challenges** On a bump-in-the-wire encryption setup, an attacker intercepting traffic between the bump in the wire will only see encrypted traffic across any potential attacker-controlled parts of the network, preventing an attacker from reading or modifying the traffic passing through, although this challenge was not encountered due to mis-configuration. Upon obtaining the SSH keys from the password-less SSH tunnel hosts, decryption of the tunnel traffic was investigated. Closer inspection of captured packets revealed the SSH handshake negotiating Diffie-Hellman key exchange, passing randomly generated session values for calculation of a shared secret. Also seen in the packet inspection was the agreed upon encryption algorithm of ChaCha20. Thus, the ephemeral traffic key needed for the attacker to decrypt with the ChaCha20 algorithm was not trivially obtainable and was not further pursued.

4.2.7 Segmented Network Topology without Encryption

- **Observations** The Red Team was provided two access points, one on the ISP router's subnet (outsider access) which was bereft of inverters and the other access was on one of the subnets with a random percentage of inverters. From the outsider access, MITM was unsuccessful because there were no hosts susceptible to an ARP poisoning attack. Further attempts to pivot and deploy MITM tools were unsuccessful due to Linux package dependencies on an air-gapped network. MITM was only successful on the subnet the attacker was on. However, from both access positions, DoS and packet replay to the inverters were successful.

- **Challenges** From an outsider position on an emulated network, it is not a target-rich environment. Pivoting into subnets with targets is difficult when hosts do not have the human element and the OS vulnerabilities seen in the real-world.

4.2.8 Segmented Network Topology with Encryption

- **Observations** The addition of encryption from the segmented unencrypted network added bump-in-the-wire SSH gateway hosts on a subnet basis. Reconnaisance confirmed that the encryption tunnel was again mis-configured, with the inverters immediately connected to the ISP router rather than being located behind the SSH box. The architecture in Figure 2.11 shows exactly this mistake. While MITM was still an available attack when on the same subnet, an outsider without the ability to pivot and deploy tools remains excluded from this attack vector.

- **Challenges** No unique challenges were introduced in this topology.

4.2.9 Segmented Network Topology with HIL and without Encryption

- **Observations** The addition of a physical inverter in the topology provided a target on which the Red Team previously conducted an assessment. In that assessment, the team conducted successful reconnaissance, vulnerability scans, packet replay, MITM, sniffed passwords, DoS attack, and evaluated the bookkeeping (logs) of the device during a security event. In contrast with the previous assessment, the DER device was not on the same subnet, and thus the vendor software and DER Connection Assistant tools were unable to discover the device. SunSpec Dashboard was able to successfully connect to the DER device, and showed some minor differences from the emulated inverters in the topology–most notably the lack of a connection register previously used to disconnect the inverter's communications. DoS and packet replay were successful.

  It was shown that flipping the DER Volt-VAR curve about the reactive power-axis caused the device to sink power when phase voltages were low. This is not a desired operating mode for the equipment because it will drive the power system away from nominal voltage. The effect can be seen in Figure 4.12.

  To customize grid settings on the physical inverter, the vendor provides the user a grid guard code to be able to change the grid parameters. Simply Modbus, a non-SunSpec Modbus complicance tool was used to initiate grid guard viewing and control. The tool was able to succesfuly write to grid guard protected control values. However it was unsuccesful reading the some of the values.

  Except for the addition of the HIL, the attacks conducted on the topology were the same for the Segmented-Network Topology without Encryption.

- **Challenges** The The HIL inverter was known to have complete inverter control models and communicate with UDP. However, while attached to the Emulytics environment, the HIL inverter could not be commanded with Netcat UDP packets and the Red Team did not discover whether this was due to the Emulytics platform translating all traffic through protocol buffers or due to other network effects. Python UDP communications still succeeded. Other challenges are the same as presented in Section 4.2.7

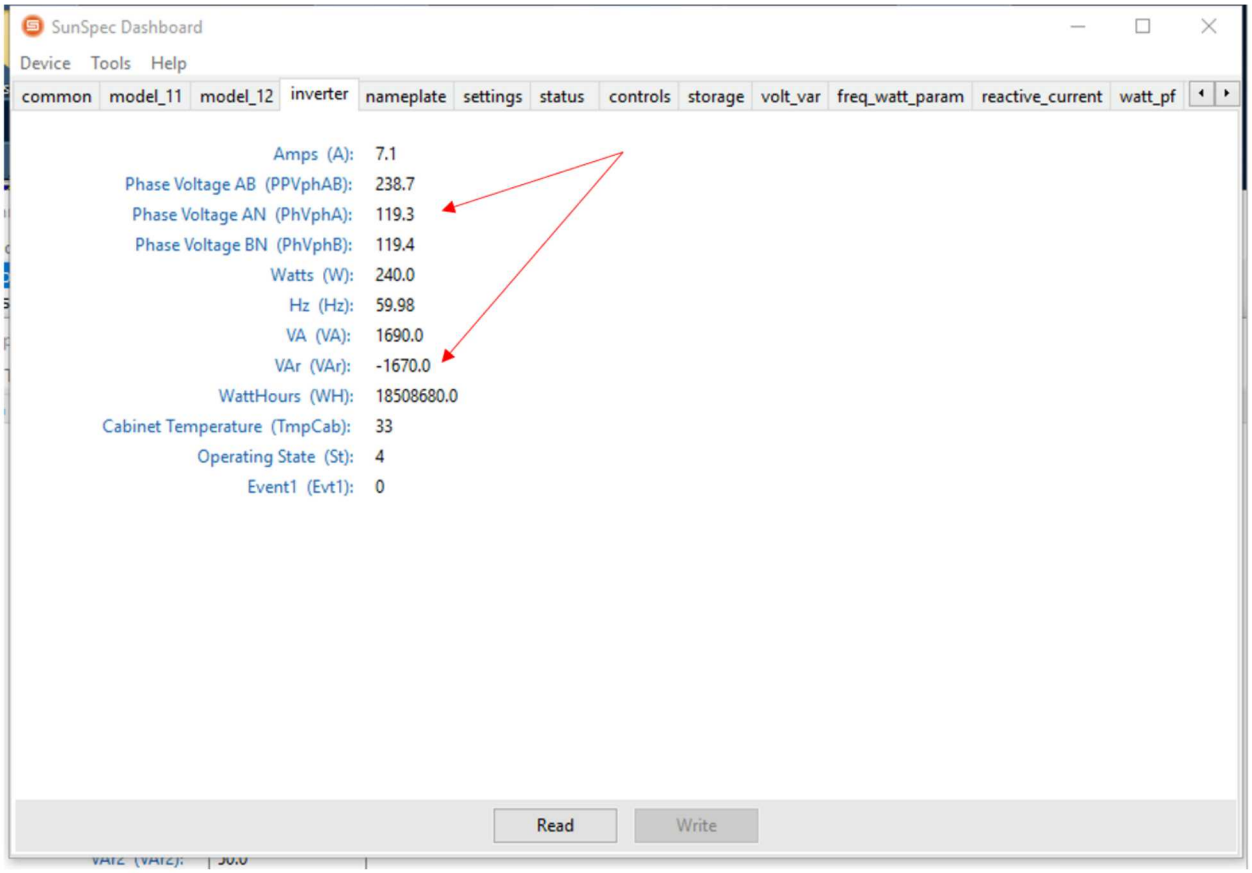

Figure 4.12. SunSpec Dashboard reading DER with inverted VV curve

4.2.10 Moving Target Defense Network Topology

- **Observations** The Moving Target Defense (MTD) environment may be a difficult toplogy to conduct reconnaissance because the networking stack implemented an IP-MAC-Port whitelisting that prevented network visibility of the DERs and the IP addresses of the equipment regularly changed. However, a security weakness - vulnerable default switch proprietary communication protocol - was created through ineffective deployment. The Red Team was able to exploit the default configurations on the switch connected to the ISP router to perform a VLAN hopping attack. This attack enabled the Red Team to listen to all broadcasts on the VLANs to gain reconnaissance information - VLAN information, IP addresses used by the SDN controller, and open ports. DoS attacks on the switch was also sucessful in preventing traffic between the utility and the DER devices. MITM attack was not successful because of the size of the IP address space that needed to be scanned for valid addresses.

- **Challenges** The MTD environment was built out with software defined networking (SDN) concepts inside of an Emulytics platform itself built on rapid prototyping models of SDN,

causing a fusion of certain network surfaces that would have been separated in the real world. For instance, a real MTD system would protect the applications and application plane communications with the interceding control plane, leaving the controller and control plane communications as a new attack surface. Conflation of the Emulytics platform and the MTD environment may have contributed to difficulties defining what elements were in scope and what new attack surfaces were available. Without the identified security weakness (which can exist in real networks), this virtual environment was far more challenging to craft MITM because the target's IP address kept changing. Access control using network function virtualization in software defined networking adds additional challenges to conducting reconnaissance on a network. However, the MTD topology did not withstand many of the attempts at reconnaissance, denial of service, packet replay, man-in-the-middle, or VLAN hopping. These attempts were prone to causing system failure, which was attributed to the novelty of the integration of the complex co-simulation sub-systems.

Finally, the common observation and challenge evident in all the topologies was the limitation from implementing an abbreviated set of registers on the testbed. This artificially limited the attack surface of the simulated inverters.

4.2.11 Summary

In theory, adding each of the cybersecurity features should improve the security posture of the DER network. As shown in Table 4.1, adding segmentation prevents adversaries outside the subnet from accessing the devices and adversaries with access to DER subnets from reaching into other enclaves. Encryption prevents replay and MITM attacks because the adversary cannot authenticate the connection to the DERMS or DER. Moving Target Defense further challenges the adversary because they cannot identify DER IP address, ports, or protocols. Denial of Service attacks are very difficult to defend against, but whitelisting the DERMS and DER can help prevents these attacks. As shown in the Table 4.1, theoretical risk scores were then calculated for Confidentiality based on the replay and MITM attacks, Integrity based on the replay and MITM attacks, and Availability based on the DoS attack.

For the CIA triad columns, a scale of 1 to 5 was created in order to categorize the risk level on each topology. A score of 1 indicates a low risk to all devices (green color code), whereas a score of 5 (red color code) indicates a high risk to a majority of the devices. Risk scores between 2 (light green color code), 3 (yellow color code), and 4 (orange color code) indicates the varying levels showing the progressive difficulty in trying to compromise some but not all the devices or the effort required to effect more devices is too great for the risk level to be ranked high.

To determine the total score, the following vulnerabililty level metrics were loosely adapted from the NIST CVSS v2.0 ratings:

- HIGH - means that means that an attack has fully succeeded. For this metric, a range of values between 10-15 is assigned.

- MEDIUM - means that attacks have partly succeeded. For this metric, a range of values between 5-9 is assigned.
- LOW - means that attacks have not succeeded. For this metric, a range of values between 0-4 is assigned.

The scores for the theoretical security were totaled for a security risk score between 3-15 from the potential score range of 0 -15. In this defined range, low risk scores between 0-4 have a green color code, medium risk scores between 5-9 have an orange color code, and high risk scores between 10-15 have a red color code.

After the red team assessments, the actual scores for each of the topologies were much different than anticpated. As shown in Table 4.2, the Red Team was successful in subverting many of the scenarios. The use of encrypted tunnels between the utility and the DERs introduced a pivot point for the attacker because of the tunnel location. The bump-in-the-wire SSH implementation did not have a password and this error was also exploited. Ultimately, the tunnel location misconfiguration exposed all the subnets to adversary control because they could directly communicate to the DER equipment in cleartext. While this was not intentional from the development team, it is realistic of deployed networks and examples of these mistakes are not uncommon ”in the wild.” This is an important result, as it reinforces the risks associated with poor network management practices. The Moving Target Defense environment had avoidable layer 2 unsecured default configurations that were exploited.

In summary, a flat topology lends itself to the attacker having layer 2 access to all devices on the network, and thus full access to the network traffic, affecting all aspects: Confidentiality, Integrity, and Availability. All attacks demonstrated in this assessment were able to be conducted: DoS, Packet Replay, and MITM. Adding encryption to the flat topology lends well to preventing packets in-transit from being read or modified. This is assuming that the attacker cannot access unencrypted traffic between the DER and the encryption point (bump-in-the-wire); however, in this assessment, the subnet of the attacker enabled visibility of the DERs and the plaintext traffic between the tunnel endpoint and the DERs, enabling all attacks just as the topology without encryption.

Segmenting the network and dispersing the DERs on separate networks removes all visibility of packets from the attacker. By this technique alone, an outside attacker is unable to read or modify packets in flight, preserving integrity and confidentiality. By assuming that network segments are protected by a firewall implementing even the simplest NAT policies, the DERs are not visible or reachable by an attacker outside the network segment, and thus packet replay is not a viable attack. In this assessment, the network segments were not protected, exposing the DERs to replay attacks. DoS attacks remain viable with single-path topologies such as a star, as the central router can be flooded. By adding encryption to the segmented topology, DERs are protected from packet replay and MITM attacks from a position of an insider.

Moving Target Defense provided a couple of features that initially inhibited red team traction. The use of SDN allowed on-switch access control. Packets not matching the whitelist for the expected IP and MAC addresses on a particular switch port were not transmitted by the switch.

Topology	Encryption	Access	Attacks			Risk Level			Total Score
			DoS	Replay	MITM	C	I	A	
Flat	None	Insider	✓	✓	✓	5	5	5	15
Flat	None	Outsider	✓	✓	✓	5	5	5	15
Flat	RFC 7539	Insider	✓			1	1	5	7
Flat	RFC 7539	Outsider	✓			1	1	5	7
Segmented	None	Insider	✓	o	o	3	3	4	10
Segmented	None	Outsider	✓			2	2	3	7
Segmented	RFC 7539	Insider	✓			1	1	4	6
Segmented	RFC 7539	Outsider	✓			1	1	3	5
Flat MTD	None	Insider	✓			1	1	5	7
Flat MTD + WL	RFC 7539	Outsider				1	1	2	4
Seg MTD + WL	RFC 7539	Outsider				1	1	2	4

- ✓ indicates the attack is possible for all DER devices
- o indicates the attack could succeed for a portion of the DER devices
- WL indicates whitelisting of the MTD network
- RFC 7539 is the IETF Protocol for the ChaCha20 stream cipher and Poly1305 authenticator

Table 4.1. Theoretical security scores for different DER communication networks.

This gave the stance of the attacker no visibility to any devices or traffic on the network besides the gateway router. This advantage was reduced when the Red Team exploited layer 2 vulnerable default configurations which made the network susceptible to some reconnaisance and DoS attacks used in disrupting communication paths.

Based on the red teaming experiments, the following are noted:

- Denial of service is difficult to prevent. Aggregators/utilities should implement firewall whitelists to prevent these types of attacks.
- Segmentation makes it difficult for the adversary to move between subnets. Flaws in system configuraion and networking implementation enabled the Red Team to manipulate all DER devices.
- Implementing the right encryption tunnel between the DERMS and DER drastically reduces the risk of Replay and MITM attacks.
- It is important that developers add layers of defense by reviewing and pushing secure code to applications.
- MTD has the potential to drastically improve security for DER networks, but this is still an area of research.

Topology	Encryption	Access	Attacks			Risk Level			Total Score
			DoS	Replay	MITM	C	I	A	
Flat	None	Insider	✓	✓	✓	5	5	5	15
Flat	None	Outsider	✓	✓	✓	5	5	5	15
Flat	RFC 7539	Insider	✓	✓	✓	5	5	5	15
Flat	RFC 7539	Outsider	✓	✓	✓	5	5	5	15
Segmented	None	Insider	✓	✓	✓	5	5	5	15
Segmented	None	Outsider	✓	✓		5	5	5	15
Segmented + PHIL	None	Outsider	✓	✓		5	5	5	15
Segmented	RFC 7539	Insider	✓	✓	✓	5	5	5	15
Segmented	RFC 7539	Outsider	✓	✓	o	5	5	5	15
Flat MTD + WL	None	Insider	✓			1	1	5	7

- ✓ indicates the attack is possible for all DER devices
- o indicates the attack could succeed for a portion of the DER devices
- WL indicates whitelisting of the MTD network
- RFC 7539 is the IETF Protocol for the ChaCha20 stream cipher and Poly1305 authenticator

Table 4.2. Security scores for different DER communication networks based on red team assessments.

This page intentionally left blank.

Chapter 5

Conclusion

A number of experiments were conducted to better understand the tradeoffs between cybersecurity features and power system performance. As new security components are added to power systems communications, there is a risk the added latency, dropouts, and reduced equipment availability could prevent effective operations. Prior research determined the acceptable latency tolerance to multiple DER communications-based distribution and transmission grid-support services. In general, the distribution schemes operated effectively with latencies of 20 seconds and larger, whereas the bulk system services (i.e., synthetic inertia, fast contingency reserves, and communication-enabled droop control) were far more sensitive–these services experienced performance issues with latencies as low as 110 ms, with certain control parameters. This work investigated the impact to DER communication rates and power system performance when adding network segmentation, SSH encryption, and MTD. It was found that these technologies add relatively small delays to the communications system ($<$10 ms) in comparison to the latencies from geographic distance (50-100 ms) or DER read/write times (2-166 ms averages).

Security improvements from each of these defensive measures were expected to be substantial, however, several network element and system configuration errors lessened the potential protections against attacks. A skilled red team conducted adversary-based assessments of multiple DER networks to quantify the improvements in DER network security from each technology. Adding DER enclaves increased defenses against widespread attacks. Encryption which could have increased the time for an attacker to gain control, was bypassed due to endpoint placement. MTD was highly successful in thwarting communication with the inverters until it was compromised.

Based on these experiments, the marginal decrease in communications speed and bandwidth are justified to significantly increase the security posture of OT networks. It is recommended that DER communication networks are segmented and encryption. Moving target defense is more challenging to implement in the field because of the required out-of-band communication network. Future work should be conducted on this promising technology to determine if possible deployment on DER networks is financially and technically practical. Additionally, the use of SCEPTRE to virtualize DER equipment, communication networks, and power systems was found to be slightly effective, but at this stage of Minmega developer expertise, it is not highly effective at comparing security methodologies in terms of QoS performance and resilience to cyberattacks. Continued development and use of the co-simulation platform is recommended to (a) assess security features for DER and other communication networks and (b) evaluate the impact of these technologies on communications and control system performance metrics.

This page intentionally left blank.

References

[1] Idaho National Laboratory. Cyber threat and vulnerability analysis of the U.S. electric sector, mission support center analysis report. Technical Report INL/EXT-16-40692, Idaho National Laboratory, August 2016.

[2] R. Smith. How a U.S. utility got hacked. *The Wall Street Journal*, December 2016.

[3] R. Smith. Russian hackers reach U.S. utility control rooms, homeland security officials say. *The Wall Street Journal*, June 2018.

[4] P. Caine. Russian-backed hackers infiltrating US power grid. *WTTW*, August 2018.

[5] D. Q. Wilber. Russian malware found on vermont electric utility laptop. *Los Angeles Times*, January 2017.

[6] D. Bradbury. Staff dust off their typewriters after malware attack. *Naked Security*, August 2018.

[7] D. Kirkpatrick. British cybersecurity chief warns of russian hacking. *The New York Times*, November 2017.

[8] K. Zetter. Inside the cunning, unprecedented hack of ukraine's power grid. *Wired*, March 2016.

[9] A. Greenberg. How an entire nation became russia's test lab for cyberwar. *Wired*, June 2017.

[10] A. Greenberg. 'Crash Override': The malware that took down a power grid. *Wired*, June 2017.

[11] R. J. Campbell. Electric grid cybersecurity. *Congressional Research Service*, September 2018.

[12] B. Gruley M. Riley, J. Dlouhy. Russians are suspects in nuclear site hackings, sources say. *Bloomberg*, July 2017.

[13] N. Perlroth. Hackers are targeting nuclear facilities, homeland security dept. and F.B.I. say. *The New York Times*, July 2017.

[14] N. Perlroth. Cyberattacks put russian fingers on the switch at power plants, U.S. says. *The New York Times*, March 2018.

[15] A. Greenberg. The untold story of notpetya, the most devastating cyberattack in history. *Wired*, August 2018.

[16] E-ISAC. Analysis of the cyber attack on the ukrainian power grid: Defense use case. Technical report, E-ISAC, March 2016.

[17] Solar Energy Industries Association and GTM Research. U.s. solar market insight q3 2018. Technical report, SEIA, September 2018.

[18] U.S. Energy Information Administration. U.S. battery storage market trends. Technical report, SEIA, May 2018.

[19] S. Page. Hawaii will soon get all of its electricity from renewable sources. *Think Progress*, May 2018.

[20] I. Penn. California invested heavily in solar power. now there's so much that other states are sometimes paid to take it. *LA Times*, June 2017.

[21] IEEE Std. 1547-2018. IEEE standard for interconnection and interoperability of distributed energy resources with associated electric power systems interfaces. Technical report, Institute of Electrical and Electronics Engineers, Inc., Feb 2018.

[22] Pacific Gas and Electric Co. Electric rule no. 21, generating facility interconnections. Technical report, Filed with the CPUC, June 2017.

[23] IEEE Std. 2030.5-2013. IEEE adoption of smart energy profile 2.0 application protocol standard. Technical report, Institute of Electrical and Electronics Engineers, Inc., Nov 2013.

[24] SunSpec Alliance. Common smart inverter profile: Ieee 2030.5 implementation guide for smart inverters, version 2. Technical report, SunSpec Alliance, March 2018.

[25] North American Electric Reliability Corporation. Critical infrastructure protection standards. https://www.nerc.com/pa/Stand/Pages/CIPStandards.aspx. Accessed: 11-14-2018.

[26] J. Johnson P. Cordeiro, J. Obert. Recommendations for trust and encryption in der interoperability standards. Technical report, Sandia National Laboratories, December 2018.

[27] C. Lai, N. Jacobs, S. Hossain-McKenzie, C. Carter, P. Cordeiro, I. Onunkwo, and J. Johnson. Cyber security primer for der vendors, aggregators, and grid operators. Technical Report SAND2017-13113, Sandia National Laboratories, December 2017.

[28] SunSpec Alliance. Sunspec der cybersecurity workgroup. https://sunspec.org/sunspec-cybersecurity-workgroup. Accessed: 11-14-2018.

[29] C. Carte D. Saleem. Certification procedures for data and communication security of distributed energy resources. Technical report, NREL Technical Report, December 2018.

[30] J. Johnson. Roadmap for photovoltaic cyber security. Technical Report SAND2017-13262, Sandia National Laboratories, December 2017.

[31] J. E. Quiroz, M. J. Reno, O. Lavrova, and R. H. Byrne. Communication requirements for hierarchical control of volt-var function for steady-state voltage. In *IEEE ISGT, Arlington, VA*, April 2017.

[32] M. Reno, J. Quiroz, O. Lavrova, and R. Byrne. Evaluation of communication requirements for voltage regulation control with advanced inverters. In *IEEE North American Power Symposium, Denver, CO*, September 2016.

[33] M. J. Reno and K. Coogan. Grid integrated distributed pv v) version 2. Technical Report SAND2014-20141, Sandia National Laboratories, 2014.

[34] Minimega. A distributed vm management tool. minimega.org. Accessed: 11-14-2018.

[35] SunSpec Alliance. Specifications & information models. https://sunspec.org/about-sunspec-specifications. Accessed: 11-20-2018.

[36] J. Stamp,C. Veitch, J. Henry, et al. Microgrid cyber security reference architecture (v2). Technical Report SAND2015-9711, Sandia National Laboratories, November 2015.

[37] J. Stamp, et al. Design tradeoffs and cyber security for microgrids. In *Energy Exchange: Federal Sustainability for the Next Decade*, Aug 2016.

[38] J. Johnson, et al. Design and evaluation of a secure virtual power plant. Technical Report SAND2017-10177, Sandia National Laboratories, 2017.

[39] OpenSSH. Openssh 7.9 release notes. http://www.openssh.com/txt/release-7.9. Accessed: 10-19-2018.

[40] A. R. Chavez, J. R. Hamlet, and W.M.S. Stout. Artificial diversity and defense security (addsec) final report. Technical Report SAND2018-4545, Sandia National Laboratories, April 2018.

[41] A. R. Chavez, W. M. S. Stout, and S. Peisert. Techniques for the dynamic randomization of network attributes. In *2015 International Carnahan Conference on Security Technology (ICCST), Taipei*, 2015.

[42] S. Hossain-McKenzie, C. Lai, A.R. Chavez, and E. Vugrin. Performance-based cyber resilience metrics: An applied demonstration toward moving target defense. In *44th IECON, Washington DC*, 2018.

[43] SunSpec Alliance. Sunspec system validation platform (svp). https://sunspec.org/sunspec-svp. Accessed: 11-14-2018.

[44] S.T. Jones, K.G. Gabert, and T.D. Tarman. Evaluating emulation-based models of distributed computing systems. Technical Report SAND2017-10634, Sandia National Laboratories, August 2017.

[45] US DOE OE. Modern distribution grid: Decision guide volume iii. Technical report, US DOE, June 2017.

[46] R. Concepcion, F. Wilches-Bernal, and R. Byrne. Effects of communication latency and availability on synthetic inertia. In *2017 IEEE ISGT, Washington DC*, April 2017.

[47] F. Wilches-Bernal, R. Concepcion, J. Neely, R. Byrne, and A. Ellis. Communication enabled – fast acting imbalance reserve (ce-fair). *IEEE Trans. Power Systems*, 33(1):201–213, January 2018.

[48] F. Wilches-Bernal, et al. Impact of communication latencies and availability on droop-implemented primary frequency regulation. In *49th NAPS, Morgantown, WV*, 2017.

[49] J. Johnson, R. Ablinger, R. Bruendlinger, B. Fox, and J. Flicker. Design and evaluation of sunspec-compliant smart grid controller with an automated hardware-in-the-loop testbed. *Technology and Economics of Smart Grids and Sustainable Energy*, 2(16), December 2017.

[50] NIST. Common vulnerability scoring system calculator version 3. https://nvd.nist.gov/vuln-metrics/cvss/v3-calculator. Accessed: 11-14-2018.

DISTRIBUTION:

1 Guohui Yuan
U.S. Department of Energy
1000 Independence Avenue SW
Washington, DC 20585

1 Kemal Celik
U.S. Department of Energy
1000 Independence Avenue SW
Washington, DC 20585

1 Dan Ton
U.S. Department of Energy
1000 Independence Avenue SW
Washington, DC 20585

1 Carol Hawk
U.S. Department of Energy
1000 Independence Avenue SW
Washington, DC 20585

1 Lee Slezak
U.S. Department of Energy
1000 Independence Avenue SW
Washington, DC 20585

1	MS 1033	Jay Johnson, 8812
1	MS 1033	Abraham Ellis, 8812
1	MS 1033	Jimmy Quiroz, 8812
1	MS 1033	Birk Jones, 8812
1	MS 0933	Brian Gaines, 09366
1	MS 0671	Jordan Henry, 05828
1	MS 0671	Jason Stamp, 05823
1	MS 0671	Derek Hart, 05821
1	MS 0671	Jennifer Depoy, 05828
1	MS 0161	Legal Technology Transfer Center, 11500
1	MS 0899	Technical Library, 9536 (electronic copy)

This page intentionally left blank.


Sandia National Laboratories