

FINAL REPORT

Critical Energy Infrastructure Cyber Defense-in-Depth

ESTCP Project EW-201607

AUGUST 2018

Kevin B. Jordan
Eamon Jordan
Kevin Ruddell
Resurgo, LLC

Distribution Statement A
This document has been cleared for public release



Page Intentionally Left Blank

This report was prepared under contract to the Department of Defense Environmental Security Technology Certification Program (ESTCP). The publication of this report does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official policy or position of the Department of Defense. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the Department of Defense.

Page Intentionally Left Blank

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 16-08-2018		2. REPORT TYPE ESTCP Final Report		3. DATES COVERED (From - To) 17 Aug 2016 – 16 Aug 2018	
4. TITLE AND SUBTITLE Critical Energy Infrastructure Cyber Defense-in-Depth				5a. CONTRACT NUMBER W912HQ-16-C-0046	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Kevin Jordan, Eamon Jordan, Kevin Ruddell				5d. PROJECT NUMBER EW-201607	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Resurgo, LLC 2800 Woodlawn Drive, Suite 141, Honolulu, HI, 96822				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Environmental Security Technology Certification Program 4800 Mark Center Drive, Suite 16F16 Alexandria, VA, 22350-3605				10. SPONSOR/MONITOR'S ACRONYM(S) ESTCP	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) EW-201607	
12. DISTRIBUTION / AVAILABILITY STATEMENT Distribution A; unlimited public release					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This project implemented, and successfully demonstrated in a power plant, an aware, fault and intrusion tolerant cyber defense-in-depth for a power plant's control system network. Two technologies; Spire, a suite of fault and intrusion tolerant technologies, provided by Johns Hopkins University and Spread Concepts, LLC, and the MANA™, Network Intrusion Detection System (NIDS) with machine-learning tools for cyber situational awareness, provided by Resurgo, LLC, and aligned with the Spire technologies, were demonstrated during the operational validation in the power plant. The MANA™ NIDS was able to detect all attacks from initial reconnaissance through exploit and port denial of service to hide the exploits. All systems protected by Spire were not compromised. Additionally, neither technology induced latency nor errors in plant control systems, communications, or devices. Costs, assuming all new equipment and installation by the technology providers, for one installation is estimated to be \$250,000.00, not including travel					
15. SUBJECT TERMS Artificial Intelligence, Machine-learning, Network Intrusion Detection					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Kevin Ruddell
a. REPORT UNCLASS	b. ABSTRACT UNCLASS	c. THIS PAGE UNCLASS			19b. TELEPHONE NUMBER (include area code) 808-784-0562
			UNCLASS	115	

Page Intentionally Left Blank

FINAL REPORT

Project: EW-201607

TABLE OF CONTENTS

	Page
EXECUTIVE SUMMARY	ES-1
1.0 INTRODUCTION.....	11
1.1 BACKGROUND.....	12
1.2 THE TECHNOLOGIES AND THEIR BENEFITS.....	12
1.2.1 Phase 1: Demonstration at PNNL	13
1.2.2 Phase 2: Demonstration at HECO Honolulu Plant.....	15
1.3 OBJECTIVE OF THE PROJECT.....	16
1.4 REGULATORY DRIVERS.....	18
2.0 TECHNOLOGY DESCRIPTION.....	21
2.1 TECHNOLOGY OVERVIEW	21
2.1.1 Operations Network Defense	21
2.1.2 Cyber Situational Awareness	26
2.2 TECHNOLOGY DEVELOPMENT	28
2.2.1 Operations Network Defense Technology Development.....	28
2.2.2 Cyber Situational Awareness Technology Development.....	33
2.3 ADVANTAGES AND LIMITATIONS OF THE TECHNOLOGY	36
3.0 PERFORMANCE OBJECTIVES.....	39
3.1 PHASE 1 - PNNL RISK REDUCTION (LABORATORY)	39
3.1.1 Phase 1 - PNNL Risk Reduction (Laboratory) Performance Objectives	39
3.1.2 Phase 1 - PNNL Risk Reduction (Laboratory) Performance Objectives Descriptions.....	39
3.2 PHASE 2 - HECO OPERATIONAL VALIDATION (Power Plant)	42
3.2.1 Phase 2 - HECO Operational Validation (Power Plant) Performance Objectives.....	42
3.2.2 Performance Objectives Descriptions	43
4.0 FACILITY/SITE DESCRIPTION	45
4.1 FACILITY/SITE LOCATION AND OPERATIONS	45
4.1.1 Pacific Northwest National Laboratory	45

TABLE OF CONTENTS (Continued)

	Page
4.1.2 Hawaiian Electric Company Power Plant	46
4.2 FACILITY/SITE CONDITIONS.....	46
4.2.1 PNNL.....	46
4.2.2 HECO Power Plant.....	47
5.0 TEST DESIGN.....	49
5.1 PNNL CONCEPTUAL TEST DESIGN.....	49
5.1.1 Level A: Enterprise Network	53
5.1.2 Level B: Operations Network.....	54
5.1.3 Baseline Characterization.....	56
5.1.4 Design and Layout of Technology Components.....	57
5.1.5 Operational Testing	58
5.1.6 Sampling Protocol	60
5.1.7 Sampling Results.....	60
5.2 HECO OPERATIONAL VALIDATION CONCEPTUAL DESIGN	60
5.2.1 Conceptual Test Design.....	61
5.2.2 Baseline Characterization.....	62
5.2.3 Design and Layout of Technology Components.....	62
5.2.4 Operational Testing	65
5.2.5 Sampling Protocol	67
5.2.6 Sampling Results.....	67
6.0 PERFORMANCE ASSESSMENT.....	69
6.1 Phase 1 - Risk Reduction Demonstration (PNNL) Data Analysis	69
6.1.1 Performance Objective 1 – Architecture Cyber Security	69
6.1.2 Performance Objective 2 – NIDS Performance	70
6.1.3 Performance Assessment.....	72
6.2 Phase 2 - Operational Validation (Power Plant) Data Analysis.....	77
6.2.1 Phase 2, Operational Validation (Power Plant), Performance Objectives Descriptions.....	77
6.2.2 Performance objective 1 – Technology Compatibility.....	78
6.2.3 Performance Objective 2 – NIDS Performance	78
6.2.4 Data Analysis	78
6.2.5 Performance Assessment.....	79

TABLE OF CONTENTS (Continued)

	Page
7.0 COST ASSESSMENT	81
7.1 COST MODEL.....	81
7.2 COST DRIVERS.....	83
7.3 COST ANALYSIS AND COMPARISON	84
8.0 IMPLEMENTATION ISSUES.....	87
9.0 REFERENCES.....	91
APPENDIX A POINTS OF CONTACT	A-1
APPENDIX B CALIBRATION OF EQUIPMENT.....	B-1
APPENDIX C QUALITY ASSURANCE SAMPLING	C-1

LIST OF FIGURES

	Page
Figure 1. The MANA™ Process	14
Figure 2. CTAM Data Flow	15
Figure 3. Abstraction of How Prime Defends	23
Figure 4. Approximation of CTAM Analysis of a Detected Attack	27
Figure 5. Spire Architecture with a Single Control Center	29
Figure 6. MANA™ Employment Schematic	33
Figure 7. CTAM-R Employment Schematic.....	35
Figure 8. The Pacific Northwest National Laboratory Campus	45
Figure 9. The HECO Honolulu Power Plant	46
Figure 10. Depiction of Test Environment.....	54
Figure 11. Depiction of The PNNL Test Environment	55
Figure 12. System Depiction	58
Figure 13. System Depiction with Defense-in-Depth Components	63
Figure 14. Alert Analysis Flow Chart	71
Figure 15. DELL PowerEdge R720 server	87
Figure 16. Spire and Switch in a Rack In the Power Plant (4U high).....	88

LIST OF TABLES

	Page
Table 1. All Runs Aggregate Attack Recall Scores	5
Table 2. Replica Requirements for Guaranteed Performance	25
Table 3. PNNL Risk Reduction Performance Objectives	39
Table 4. CIA Scoring Example	40
Table 5. Performance Objectives	42
Table 6. Timeline of PNNL Test Activities (March-April 2017)	59
Table 7. Timeline of the HECO Demonstration (Jan-Feb 2018)	66
Table 8. Performance Objective One	69
Table 9. CIA Scoring Example	69
Table 10. NIDS Recall Performance Objective	70
Table 11. NIDS Precision Performance Objective.....	70
Table 12. Architecture Type Aggregate Shapiro-Wilk Test	72
Table 13. Architecture Type Confidentiality Shapiro-Wilk Test.....	72
Table 14. Architecture Type Integrity Shapiro-Wilk Test	73
Table 15. Architecture Type Availability Shapiro-Wilk Test	73
Table 16. Architecture Type Basic Statistics for CIA and Aggregate Score	73
Table 17. Architecture Type Aggregate Paired T-Test	74
Table 18. Architecture Type Confidentiality Paired T-Test.....	74
Table 19. Architecture Type Integrity Paired T-Test	74
Table 20. Architecture Type Availability Paired T-Test.....	74
Table 21. Architecture Type Aggregate Wilcoxon Sign-Rank Test	75
Table 22. Red Team Access Location Aggregate Shapiro-Wilk Test.....	76
Table 23. Red Team Access Location Basic Statistics for CIA and Aggregate Score	76
Table 24. Red Team Access Location Aggregate Paired T-Test	77
Table 25. Performance Objective 1	77
Table 26. Performance Objective 2	78
Table 27. Estimated Cost Model for the Defense-in-Depth	81

Page Intentionally Left Blank

ACRONYMS AND ABBREVIATIONS

ANOVA	Analysis of Variance
API	Application Program Interface
APT	Advanced Persistent Threat
C2	Command and Control
C4I	Command, Control, Communications, Computers and Intelligence
CD	Compact Disc
CEH	Certified Ethical Hacker
CIA	Confidentiality, Integrity and Availability
CIP	Critical Infrastructure Protection
CMP	Cyber Monitoring Plan
COTS	Commercial Off-the-shelf
CSD	Computer Science Department
CTAM-R	Cyber Threat Activity Matrix with Radiflow ISID
DAD	Data Acquisition Daemon
DARPA	Defense Advanced research Projects Agency
DHS	Department of Homeland Security
DMZ	Demilitarized Zone
DNP3	Distributed Network Protocol
DoD	Department of Defense
DoE	Department of Energy
EMS	Energy Management System (EMS)
EO	Executive Order
ESTCP	Environmental Security Technology Certification Program
EW	Energy and Water
EXWC	Expeditionary Warfare Command
FBI	Federal Bureau of Investigation
GUI	Graphical User Interface
HASP	Health and Safety Plan
HECO	Hawaiian Electric Company
HMI	Human Machine Interface
ICS	Industrial Control Systems
IDS	Intrusion Detection System
ISA	Interconnection Security Agreement
iSID	Industrial Cyber Security Intrusion Detection
ISO	International Organization for Standardization
IT	Information Technology (the business network for a power plant)

JHU	Johns Hopkins University
JIOR	Joint Information Operations Range
KW	Kilowatt
LLVM	Low Level Virtual Machine
MANA™	Machine-learning Assisted Network Analyzer
MB	Megabyte
MIT-LL	Massachusetts Institute of Technology – Lincoln Laboratory
MitM	Man-in-the-Middle
MLS	Machine-Learning Sensor
MW	Megawatt
NAVFAC	Naval Facilities Engineering Command
NCCIC	National Cybersecurity and Communications Integration Center
NERC	North American Electric Reliability Corporation
NIDS	Network Intrusion Detection System
NIST	National Institute of Standards and Technology
NMAP	Network Mapper
NORTHCOM	United States Northern Command
O&M	Operations and Maintenance
OT	Operational Technology (the control system network managing the PLCs and RTUs)
PACOM	United States Pacific Command
PERA	Purdue Enterprise Reference Architecture
PIT	Platform Integrated Technology
PLC	Programmable Logic Controller
PNNL	Pacific Northwest National Laboratory
POC	Point of Contact
RBS	Red Balloon Security
RIF	Rapid Innovation Fund
RRE	Risk Reduction Event
RTU	Remote Terminal Unit
SA	Situational Awareness
SCADA	Supervisory Control and Data Acquisition
SIEM	Security Information and Event Management (device)
SME	Subject Matter Expert
SNL	Sandia National Laboratory
SPI	Smart Packet Inspection
STIX	Structured Threat Information eXpression

TAXII	Trusted Automated eXchange of Indicator Information
TB	Terabyte
TBD	To Be Determined
TRL	Transition Readiness Level
TTP	Tactics, Techniques and Procedures
UFC	Unified Facilities Criteria
UPS	Universal Power Supplies
USAFA	US Air Force Academy
VLAN	Virtual Local Area Network (LAN)
VPN	Virtual Private Network

ACKNOWLEDGMENTS



Resurgo, LLC wishes to formally acknowledge the assistance and hard work of (in alphabetical order):

- The Hawaiian Electric Company (HECO), their management, information assurance department and the plant engineers who ensured we could operate.
- The Naval Facilities Engineering Command, Headquarters, Washington, D.C. (NAVFAC), and the Naval Facilities Engineering Command - Expeditionary Warfare Command (NAVFAC - EXWC)
- Pacific Northwest National Laboratories (PNNL)
- Sandia National Laboratory Cyber Red Team
- U.S. Northern Command (NORTHCOM)
- U.S. Pacific Command (PACOM)

And our technology providers:



- Imprimis, Inc. (CTAM)



- Johns Hopkins University Computer Science Department
- Spread Concepts, LLC

ABSTRACT

INTRODUCTION

The Department of Defense (DoD) and commercial industry lack the ability to effectively defend the control system networks of critical energy infrastructure against cyber-attacks. To address this capability gap, Resurgo, LLC, partnered with: Johns Hopkins University (JHU), Spread Concepts LLC, Pacific Northwest National Laboratory (PNNL), and the Hawaiian Electric Company (HECO), to resolve this National need.

OBJECTIVE

To implement, and successfully demonstrate in a power plant, an aware, fault and intrusion tolerant cyber defense-in-depth for a power plant's control system network.

TECHNOLOGY DESCRIPTION

Of the five technologies selected, only two technologies, MANA™ (Machine-learning Assisted Network Analyzer) and Spire suite were used in our operational validation to successfully defend the HECO power plant. Spire defended the network and certain hardware and the MANA™ Network Intrusion Detection System (NIDS) provided the operator cyber situational awareness (SA) missing from the Spire technologies.

PERFORMANCE AND COST ASSESSMENT

Performance Assessment

The MANA™ NIDS was able to detect all attacks from initial reconnaissance through exploit and port denial of service to hide the exploits. As a passive NIDS, MANA™ did NOT induce latency or errors in the plant's control systems, communications, or devices.

No systems protected by Spire were compromised. Illegal commands given to Spire protected plant equipment were ignored by the multi-compiled Prime replications within Spire. Additionally, neither MANA™ nor Spire induced latency or errors in the plant's control systems, communications, or devices. Plant engineers observed that in some areas Spire's response time was faster (approximately twice as fast) than the baseline system.

Cost Assessment

Costs for acquiring, installing and operating a cyber aware, fault and intrusion tolerant defense vary since the entire defense can be installed by the user / purchaser of the software and equipment. Assuming all new equipment, and installation by the technology providers, total cost for one installation is estimated to be \$250,000.00, not including travel.

IMPLEMENTATION ISSUES

The primary implementation issue will be resistance to change. The cyber defense equipment and procedures are standard for IT (Information Technology) networks but will be new inside a power plant on operational technology (OT) networks. Plant personnel will only have one additional screen to monitor, but it will provide cyber situational awareness of threats to the plant.

PUBLICATIONS

No publications, beyond Environmental Security Technology Certification Program (ESTCP) required reports, were produced by Resurgo, LLC. However, Johns Hopkins University has published numerous articles relating to its technologies.

EXECUTIVE SUMMARY

INTRODUCTION

Resurgo, LLC, is a woman-owned, small business of 13 employees in the cybersecurity sector. Years of working with machine-learners in DoD projects, led us to an understanding of machine-learning's weaknesses and how to best harness its strengths. Machine-learners work best in a simple environment that changes very little, such as industrial control systems. But it takes more than a new type of NIDS to defend against today's threats to critical cyber infrastructure. Supervisory Control and Data Acquisition (SCADA) systems must ensure continuous availability and correct operation in the presence of compromises and attacks at both the system and network level. Unauthorized access to an OT network of a US utility company by a cyber adversary represents a worst-case scenario from a critical infrastructure protection perspective. Yet, we face this existential threat on a national scale. An active and aware cyber defense-in-depth of critical infrastructure is crucial to closing this vulnerability. We were familiar with the National Institute of Standards and Technology (NIST) discussions of defense in depth (NIST Computer Security Division, 2013) – an information security strategy integrating people, technology, and operations capabilities. This definition stops at the firewalls before the operational technology networks. To address this capability gap, Resurgo, LLC, as the Prime Contractor, partnered with Imprimis, Inc., Johns Hopkins University, and Spread Concepts, Inc., to address this National need. The ESTCP funded our solution entitled “Critical Energy Infrastructure Cyber Defense-in-Depth”. We were supported on this project by the PNNL, Sandia National Laboratory, and the HECO. This team allowed Resurgo to demonstrate the first ever successful operation of a cyber aware, fault and intrusion tolerant defense of a utility control system in a functional power plant.

THE PROBLEM

The DoD lacks the ability to effectively defend the control system networks and SCADA systems of Posts, Bases, Camps, Installations, and related civilian critical energy infrastructure against cyber-attacks. The DoD's control system networks are subject to a broad spectrum of threats, the most serious of which are those associated with nation state actors and Advanced Persistent Threats (APT). The DoD actively defends its Information Technology networks, but this active defense uses funding, manpower and focused situational awareness that does not extend to DoD's Operations Technology networks. Commercial U.S. utility companies have a similar problem to the DoD in that they do not have the tools, expertise, manpower or budget to defend their OT networks from an APT.

THE PROBLEM'S IMPACT

The DoD's cyber systems, to include its energy production control systems, are probed every day by thousands of varied sources. DoD is well defended in its IT networks. However, the DoD is unable to identify new compromises in its OT systems or to determine if any of its energy control systems are already compromised. The compromise or loss of one or more energy production facilities at the beginning of an attack on the Nation, or even on a single installation, will result in the loss of Command, Control Communications, Computers and Intelligence (C4I) at a crucial moment and could result in the loss of the opening battle of that war.

This problem is compounded by the near total lack of tools made for the multi-protocol OT networks, the limited amount of IT personnel with OT knowledge – or OT personnel with knowledge of IT defensive technologies - and no budget to acquire the tools or expertise to close this massive vulnerability in the Nation’s critical energy infrastructure. A NIST-compliant utility control system network is designed to deny and restrict cyber threats from reaching the OT networks from the Internet. However, introduction of internet protocol devices and internet connectivity into SCADA spaces, and sometimes directly to field devices, creates vulnerabilities in previously non-IP, NIST compliant areas. Once a cyber threat breaches a utility’s defenses and gains access to the OT networks, being NIST-compliant does not adequately protect the critical devices and their associated data exchanges.

OBJECTIVES

The objective of this project was to implement, and successfully demonstrate, an aware, fault and intrusion tolerant cyber defense-in-depth for a power plant’s control system network. A power plant operating with the cyber defense-in-depth technologies used in this project will be cyber-secure, Byzantine fault tolerant, and aware. A plant thus protected would be able to withstand nation state equivalent cyber-attacks. Demonstrating this cyber defense onsite in a commercial utility company’s power plant shows a practical cyber-secure path forward for DoD, Department of Homeland Security (DHS), and commercial utilities.

TECHNOLOGY DESCRIPTION

In Phase 1, five technologies were involved but only four were demonstrated; two intrusion detection systems (a machine-learning IDS and the CTAM-R (Cyber Threat Activity Matrix with Radiflow ISID)) and two network defense technologies (Spines and Prime).

Phase 1 – Pacific Northwest National Laboratories Demonstration

Five cyber defense-in-depth technologies, (Prime, Spines, Symbiote, CTAM-R, and a machine-learning intrusion detection system) went to PNNL. They were layered in depth throughout the control system network. Only four of the technologies were tested by the cyber red team.

Machine-Learning IDS

The IDS used machine-learning based tools to catch new and morphed (altered) attacks that traditional signature sensors could not catch. The IDS monitored all traffic entering and leaving the control system network and provided a crucial network layer of situational awareness to the defense-in-depth architecture. Machine-learning based systems have been proven to be particularly effective in detecting polymorphic versions of attacks on which the IDS has been trained. Machine-learning systems have even shown the ability to catch true zero days or attacks on which it has not been trained. This machine-learning based IDS was the initial prototype of the MANA™ used in Phase 2.

Cyber Threat Activity Matrix – Radiflow

The CTAM-R is a prototype combined database and advanced, intelligent intrusion detection system (IDS) for industrial control systems (ICS) that uses both signature-based & anomaly detection.

It also compiles threat information, threat Tactics, techniques and Procedures (TTP), and detectable indicators from various, open sources into a single, coherent, ICS-specific database of threats. CTAM compiles the threats it has detected and shares those confirmed detections externally using Structured Threat Information eXpression (STIX) and Trusted Automated eXchange of Indicator Information (TAXII) formats, while simultaneously receiving threat intelligence updates. It uses the compiled observables and derived knowledge of the target network to guide IDS detection operations. Its Radiflow IDS component uses a modified Snort IDS, combined with advanced smart packet inspection to enable it to use both signature-based detection for known intrusion attempts and anomaly-based detection to identify variations from a derived model of normal network behavior. This gives it the ability to detect previously unknown attacks. The CTAM-R IDS evaluated the same traffic as the machine-learning IDS but used a completely different type of inspection method.

Although CTAM provided both IDS and operator situational awareness functions, only the IDS component was tested. The situational awareness component, actions taken by humans, and any other human or cognitive based effect was not evaluated during testing.

Prime

Prime is an intrusion-tolerant replication engine that protected the SCADA master of the SCADA architecture. Prime provides component and system performance guarantees even when the system is partially compromised by a sophisticated Byzantine replication protocol. Prime supports proactive recovery, diversity, and state transfer through its multiple Prime servers. Using its multi-compiler, Prime ensures that each of its replicas presents a different attack surface (different variation) to an attacker so that no single exploit can take down the entire system. Prime servers periodically rejuvenate to clean the system from potentially undetected intrusions. Prime seeks source congruency. That is, a strong majority of the Prime servers must have the same result, or outcome, for an operation to be congruent.

Spines

The Spines overlay communications architecture ensured delivery of SCADA traffic within very tight time constraints while blocking adversary positioning for man-in-the-middle-attacks. Spines provided a highly reliable and fast messaging framework that overlaid the network structure connecting the Human Machine Interface (HMI) with the SCADA devices it manages. Speedy and reliable delivery of Spines supported traffic is accomplished by multicasting network intelligence on the status of packets and availability of open routes across the overlay. A derivative of the Spines technology is in commercial use providing reliable and timely video delivery services worldwide for media enterprises such as Fox News, CNN, NPR and Bloomberg Financial Services.

Symbiote

Symbiote completes the cyber defense-in-depth by providing host-based endpoint protection of Remote Terminal Units (RTUs), Programmable Logic Controllers (PLCs) and other devices. Symbiote detects and reverses firmware intrusions and modifications. Symbiote provides SCADA networks with critical endpoint device defense where none exists currently. Symbiote uses Automatic Binary Structure Randomization to protect the firmware by changing the structure and code.

The result of this change is that the protected embedded device is no longer vulnerable to compromise. Symbiote is source code agnostic and can be used on any device or any manufacturer.

Symbiote is the ideal host (PLC, RTU, printer, etc.) level defense to complete the defense-in-depth. Red Balloon Security, Symbiote's technology provider, withdrew from the project and did not give us the software to defend the PLCs or RTUs. Rather than not have Symbiote represented in the experiment, we purchased a commercial networked printer defended by Symbiote software and included it within the defended OT network to see if the red team could exploit it.

Phase 2 – Hawaiian Electric Company Power Plant Demonstration

From the SME interactions and PNNL results, the machine-learning IDS, Spines and Prime were determined to be strong enough to be used at HECO. The machine-learning IDS evolved into MANA™ (Machine-learning Assisted Network Analyzer) and Spines and Prime were upgraded, further hardened, and combined with other applications to become the Spire suite of technologies. In Phase 2, Spire defended the network and certain hardware and the MANA™ Network Intrusion Detection System (NIDS) provided the plant operator the cyber situational awareness (SA) missing from the Spire technologies.

Unfortunately, we lacked the funding to continue use of the CTAM NIDS. However, due to HECO security policies, the CTAM would not have been able to contribute to the demonstration since no connections external to the plant were permitted on networked equipment. Without the external connection CTAM would not have been able to demonstrate the power of its database.

The MANA™ NIDS is an evolutionary upgrade of the Phase 1 machine-learning IDS' components and processes. Different machine learning algorithms were integrated, training techniques were modified, correlation of alerts was moved internal to MANA™, a new feature was added so that a confidence weighting is assigned to each alert, and the graphical user interface was simplified and improved so that the information displayed to the human plant engineer is prioritized with links to the supporting data for further analysis. The MANA™ NIDS is passive and receives all traffic via a one-way network tap. The MANA™ NIDS is compatible with all networks using IP based traffic to include variations such as MODBUS over TCP/IP.

Spire is a suite of fault and intrusion tolerant technologies based on the successful Spines and Prime technologies, with the addition of:

- A customized and hardened HMI derivative of the open-source pvbrowser (an application framework that provides a specialized browser for the client computer and an integrated development environment for the data acquisition programs that connect to the SCADA Master).
- A hardened PLC/RTU Proxy (a device that relays Modbus or Distributed Network Protocol (DNP3) commands between the PLC or RTU and the SCADA Master).
- A SCADA Master designed from scratch to be intrusion tolerant and compatible with replication using the Prime intrusion-tolerant replication engine. The SCADA Master is the core of a typical SCADA system. It collects data from the PLCs and RTUs in the field and sends commands to control the state of the field devices.

Spire is compatible with all networks using IP based traffic to include MODBUS over TCP/IP.

Technology Notes

Fault and intrusion tolerant technologies, whether added to the network, installed alongside as complementary to network devices (as was done here), or integrated into a replacement device, are necessary to protect OT networks and SCADA devices and to ensure continuity of energy production despite existing or introduced vulnerabilities. Fault and intrusion tolerant technologies significantly raise the cost of entry for APT but are not foolproof or a silver bullet. The aware cyber Defense-in-Depth concept for utility control system defense also uses intrusion detection sensing. The NIDS is the only defensive technology that reports to the plant engineers and network defenders about what is happening on and to the control system and its networked devices.

The correlated, high confidence alerts provided by the NIDS provide details about the threats attacking the formidable barriers provided by the fault and intrusion tolerant technologies. The existence of fault and intrusion tolerant systems on a utility OT network will force the attacker to adopt more intrusive measures, or adjust their attack timeline, as they attempt to overcome the defenses. This in turn exposes the attacker to a higher probability of detection by the aligned NIDS and permits plant personnel to respond sooner and limit the penetration and potential damage.

NOTE: These technologies complement the NIST and North American Electric Reliability Corporation (NERC) - Critical Infrastructure Protection (CIP) requirements and add layers of defense where currently none exist. Both technologies are tested and shown not to induce latency or errors in plant equipment or networks.

PERFORMANCE AND COST ASSESSMENT

Performance Assessment

Phase 1 – PNNL

The designed experiment conducted in the PNNL lab produced the recall results shown in the table below. The table shows the total number of Sandia National Labs cyber red team attacks (53) and the number of attacks identified and missed by the sensors listed on the left of the table:

Table 1. All Runs Aggregate Attack Recall Scores

Sensor	Total Number of Red Team Attacks	Total Number of Attacks Caught	Total Number of Attacks Missed	Recall
MANA (Machine Learning)	53	42	11	79.25%
MANA (Anomaly)	53	35	18	66.04%
Radiflow (Anomaly)	53	15	38	28.30%
Radiflow (Signature)	53	1	52	1.89%

Recall is defined as the number of attacks caught over the total number of attacks conducted:

$$Recall = \frac{Caught\ Attacks}{Missed + Caught\ Attacks} = \frac{Caught\ Attacks}{Total\ Attacks}$$

The Phase 1 aggregate attack recall results revealed some interesting insights:

- A. MANA's machine learning detection did very well (comparatively) in catching the red team attacks. MANA's anomaly-based detection performed adequately in detecting attacks; however, the previous precision results showed that the MANATM anomaly alerts were not always accurate.
- B. Signature-based detection is not appropriate for detecting advanced threats targeting SCADA systems. Based on the uniqueness of each plant's equipment composition and legacy networks, attackers only use known exploits (attacks with known signatures) during the initial "establish a foothold" phase of the attack. Once inside the OT network, attackers must improvise and craft specifically targeted attacks to achieve their objectives.
- C. Radiflow's anomaly-based detection did not catch many of the red team attacks, although the alerts that were produced were the most accurate (as confirmed from the high precision results).

Phase 2 – HECO

The MANATM NIDS was able to detect 100% of the initial reconnaissance and all follow-on attacks that were later conducted. MANATM's success was determined by comparing the high confidence correlated alerts (NIDS output) against the cyber red team's attack logs to provide ground truth that there was an attack, the type of attack, when it was launched and against what IP address.

The cyber red team was unable to compromise any replicated device within the Spire defended section of the network. Commands given to devices to exceed their safety parameters were ignored by the Prime multicomplied replications. The cyber red team was also unable to isolate the defended devices from the control system or to perform DoS attacks on the ports that interfaced with the protected devices. This is consistent with results from the Phase 1 designed experiment at PNNL.

Additionally, data gathered by HECO showed that Spire did NOT induce latency or errors in control system's communications or devices. In fact, plant engineers observed that in some areas Spire's response time was faster (approximately twice as fast) than the baseline system.

Cost Assessment

Costs for acquiring, installing and operating a cyber aware, fault and intrusion tolerant defense vary since the entire defense can be installed by the user / purchaser of the software and equipment. However, assuming all new equipment and installation by the technology providers, total cost for one installation is estimated to be \$250,000.00, not including travel.

IMPLEMENTATION ISSUES

The primary implementation issue is inertia and resistance to change. For an information technology professional, implementation of the Critical Energy Infrastructure Cyber Defense-in-Depth will be no harder than any standard day-to-day IT task using tools and technology they are familiar with. Except, that they will installing in a power plant where they have no experience and may literally be lost. Conversely, for the plant engineer or operations technology network professional, this will be all new, even though it has been done routinely in the enterprise part of the utility's network for decades. Plant personnel will also have one more screen to monitor, but it will provide near real-time, highly correlated, cyber situational awareness of threats to the plant.

RESULTS

Overall Results

The Phase 1 results showed areas where all of the technologies needed improvement. Additionally, comments from the HECO, NAVFAC EXWC and PNNL SMEs suggested other areas for improvement in the technologies' displays and controls.

The Phase 2 operational validation reinforced the Phase 1 results that an aware cyber defense-in-depth using the Johns Hopkins' fault and intrusion tolerant technologies and the aligned MANA™ NIDS provides significantly better protection for control systems than the NIST compliant base-line architecture alone. HECO's plant network used the 18 control families of NIST SP 800-53. However, to test the cyber defense-in-depth technologies many of the regulatory access, physical, and training controls were set aside to permit a focused and timely demonstration. As in Phase 1, plant compliance with the cyber specific portions of NIST offered no effective protection from a sophisticated threat. The results of the Phase 2 operational validation establish that fault and intrusion tolerant technologies can protect a utility's OT networks from an advanced and determined cyber threat with insider access. The *Black Energy* and *StuxNet* incidents show that even isolated ICS networks are subject to penetration when threats jump across air-gaps into the OT networks.

The results of the Phase 2 operational validation in a functional commercial power plant proved that the MANA™ and Spire technologies:

- Function in a power plant.
- Are compatible with power plant equipment.
- Do not induce latency or errors in plant networks or devices.
- Are highly effective at identifying cyber penetrations and preventing their disruptive impact.

The results also strongly suggest that the DoD, the DHS, and commercial utilities should consider implementation of cyber defense-in-depth architectures using fault and intrusion tolerant technologies combined with an aligned machine-learning NIDS to protect the control systems and SCADA systems of critical infrastructure.

MANA™ Specific Results

The MANA™ NIDS performed better than anticipated, catching Network Mapper (NMAP) and reconnaissance conducted by the cyber red team, as well as the deliberate attacks done during the demonstration, with almost zero false positives and no false positives among the high confidence alerts. This was a marked improvement in MANA™ performance from the Phase 1 experiment results at PNNL.

In Phase 2, MANA™ had no false positives in the high confidence correlated alerts group. The MANA's patented correlation processes eliminated all false positives from the "high confidence" alerts forwarded to the human operator / plant defender for action. The high confidence alerts had ZERO false positives. However, lower confidence level alerts did include some false positives. We believe that refinement of the patented correlation process and other patented proprietary adjustments enabled this improvement in the NIDS' performance.

Spire Results

The Spire fault and intrusion tolerant technologies worked as intended without inducing latency, errors or failure in the control system being defended or the devices connected to the control system. The fault and intrusion tolerant technologies secured the control system's functionality, despite the presence of the cyber red team within the network. Additionally, data gathered by the plant's engineers showed that normal system latency was decreased in the areas defended by the Spire technologies. NOTE: HECO retained all network data due to confidentiality concerns about the information the data showed about plant cyber security, equipment and performance.

The cyber red team was unable to compromise any Prime replicated device within the Spire defended section of the network. Commands given to devices to exceed their safety parameters were ignored by the Prime multi-compiler replications defending the devices. The cyber red team was also unable to isolate the defended devices from the control system or to perform DoS attacks on the ports that interfaced with the protected devices. This is consistent with results from the Phase 1 designed experiment at PNNL. However, these same attacks were dramatically successful against the undefended portion of the plant network.

CONCLUSION

This operational validation showed that the two cybersecurity technologies, MANA™ and Spire, work as intended without inducing latency, errors, or failure in a functioning commercial power plant. The operational validation also confirmed that the Spire technologies, when combined with aligned machine learning sensing, provides better cybersecurity of a control system network than that of NIST-compliant architecture alone.

This project showed conclusively that the Spire fault and intrusion tolerant technologies are powerful in and of themselves. But, combined with an aligned MANA™ NIDS (using heterogeneous sensor architecture and an internal correlation engine), the cyber aware fault and intrusion tolerant defense strips the initiative from the hackers and conveys it to the defenders in a way that delivers a true "fight through" capability for critical infrastructure dependent on industrial control systems. We recommend that the DoD, government agencies and organizations, and utility companies start employing the fault and intrusion tolerant technologies

with appropriate machine learning sensing methodologies so that the Nation's Industrial Control Systems are able to operate unimpeded *when, not if*, intrusions occur. The combination of the Spire fault and intrusion tolerant capabilities with the greatly enhanced situational awareness provided by the MANATM machine learning offers the defenders of the Nation's critical infrastructure the opportunity to seize and hold the initiative in the cyber domain. These technologies complement NIST and NERC-CIP and add layers of defense where currently there are none.

Page Intentionally Left Blank

1.0 INTRODUCTION

Resurgo, LLC is a woman-owned, small business of 13 employees in the cybersecurity sector. Years of working with machine-learners in DoD projects, led us to an understanding of machine-learning's weaknesses and how to best harness its strengths. Machine-learners work best in a simple environment that changes very little, such as industrial control systems. But it takes more than a new type of NIDS to defend against today's threats to critical cyber infrastructure. SCADA systems must ensure continuous availability and correct operation in the presence of compromises and attacks at both the system and network level. Unauthorized access to an OT network of a US utility company by a cyber adversary represents a worst-case scenario from a critical infrastructure protection perspective. Yet, we face this existential threat on a national scale. An active and aware cyber defense-in-depth of critical infrastructure is crucial to closing this vulnerability. We were familiar with the National Institute of Standards and Technologies discussions of defense in depth (NIST Computer Security Division, 2013) – an information security strategy integrating people, technology, and operations capabilities. This definition stops at the firewalls before the operational technology networks. To address this capability gap, Resurgo, LLC, as the Prime Contractor, partnered with Imprimis, Inc., Johns Hopkins University, and Spread Concepts, Inc., to address this National need. The Environmental Security Technology Certification Program funded our solution entitled “Critical Energy Infrastructure Cyber Defense-in-Depth”.

The Critical Infrastructure Cyber Defense-in-Depth concept combined multiple technologies to form an inner bastion within the OT network. An encrypted communications overlay secures the backbone, fault and intrusion-tolerant replication of key devices secures the endpoints and a machine-learning network intrusion detection system provides cyber situational awareness and correlated alerts assigned a confidence level of the probability of the alert being an attack to the human operators and defenders of the power plant.

To test our defensive concept, we broke the project into two phases, Phase 1 would be rigorous laboratory testing at a National Lab and Phase 2 would be an operational validation of the best technologies from Phase 1 in a HECO power plant.

Phase 1 was a designed experiment conducted at the PNNL using a Sandia National Laboratory cyber red team to test the technologies. Sandia is renowned for its red team work in multiple fields including industrial control systems. PNNL was chosen because the emulation environment at PNNL is renowned for its ability to replicate SCADA systems. Additional considerations include that PNNL is a government laboratory and with numerous SMEs on SCADA, cyber security and red teams and that PNNL is one of several Department of Energy laboratories that comprise the National SCADA Test Bed Resurgo worked with HECO engineers and the PNNL team to design a machine-in-the-loop emulation that mirrored as closely as possible a HECO power plant's control system network. Phase 1 was also a risk reduction event. By mirroring the HECO power plant at PNNL, we allayed multiple SMEs' concerns about the technologies inducing problems with plant equipment and systems.

Phase 2 was an operational validation of the best Phase 1 technologies conducted in HECO power plant that was offline for maintenance. HECO partnered with Resurgo because HECO provides power to all DoD sites and installation on Oahu.

HECO's parent company, Hawaiian Electric Industries, through its other subsidiaries, provides power to all DoD sites and installations in the State of Hawaii.

The two most successful technologies of the five used in our Phase 1 risk reduction experiment PNNL were brought forward to the operational validation. The two technologies were the Spire fault and intrusion tolerant cyber defense technologies and the MANA™ NIDS and correlation engine. The cyber SA provided by the MANA™ NIDS complements Spire and ensures a complete cyber defense solution that permits plant personnel to quickly understand the developing threat and take appropriate response. An informed and aware defense is inherently more survivable than one that relies exclusively on the strength of the defense because awareness sets the conditions for the defenders to seize and maintain the initiative.

1.1 BACKGROUND

The Problem. The DoD lacks the ability to effectively defend the control system networks and SCADA systems of Posts, Bases, Camps, Installations, and related civilian critical energy infrastructure against cyber-attacks. The DoD's control system networks are subject to a broad spectrum of threats, the most serious of which are those associated with nation state actors and Advanced Persistent Threats. The DoD actively defends its IT networks, but this active defense uses funding, manpower and focused situational awareness that does not extend to DoD's OT networks. Commercial U.S. utility companies have a similar problem to the DoD in that they do not have the tools, expertise, manpower or budget to defend their OT networks from an APT.

The Problem's Impact. The DoD's cyber systems, to include its energy production control systems, are probed every day by thousands of varied sources. In the IT area, DoD is well defended. However, the DoD is unable to identify new compromises in its OT systems or to determine if any of its energy control systems are already compromised. The compromise or loss of one or more energy production facilities at the beginning of an attack on the Nation, or even on a single installation, will result in the loss of C4I at a crucial moment and could result in the loss of the opening battle of that war. This problem is compounded by the near total lack of tools made for the multi-protocol OT networks, the limited amount of IT personnel with OT knowledge – or OT personnel with knowledge of IT defensive technologies - and no budget to acquire the tools or expertise to close this massive vulnerability in the Nation's critical energy infrastructure. A NIST-compliant utility control system network is designed to deny and restrict cyber threats from reaching the OT networks from the Internet. However, introduction of internet protocol devices and internet connectivity into SCADA spaces, and sometimes directly to field devices, creates vulnerabilities in previously non-IP, NIST compliant areas. Once a cyber threat breaches a utility's defenses and gains access to the OT networks, being NIST-compliant does not adequately protect the critical devices and their associated data exchanges.

1.2 THE TECHNOLOGIES AND THEIR BENEFITS

Current Technology and practices. The current approach to cyber defense of critical infrastructure follows a "Code of Best Practices" regimen that seeks to keep computer control systems associated with SCADA, "air gapped" from the internet, fully patched and up-to-date with the latest protection against known malicious payloads or "malware."

The DoD and commercial utility companies are assisted in their efforts by the current cyber standards for OT architectures published by the NIST and the NERC-CIP.

The NIST and NERC-CIP policies and cyber specific best practices make it harder for a cyber threat to establish itself within the OT network, but compliance with these standards and practices does not alert energy providers to a threat, does not prevent access to field devices by the threat, nor does it protect against the destructive consequences likely to follow threat access. In time, a determined and capable adversary will eventually gain access to a power plant, regardless of policies and compliance. Fault and intrusion technologies will slow the threat and prevent damage to protected equipment and networks. However, without an aligned machine-learning NIDS to provide cyber situational awareness to the plant's engineers and network defenders, the threat in time will eventually find a way in. The NIDS cyber SA limits the time available to the threat to work clandestinely, while the fault and intrusion technologies buy additional time for the defender to take action.

NIST Special Publication 800-82, Guide to Industrial Control Systems (ICS) Security notes that air gapping systems has not worked and that with the extension of internet protocol connected devices into operational technology spaces, logical isolation is no longer a viable defensive strategy. In practice, both the DoD and commercial industry connect systems to the Internet that were originally designed to, and should be, air gapped. These connections are typically made to facilitate remote monitoring and maintenance by the system's engineers. On two occasions over the past few years Resurgo has discovered wireless routers connected to supposedly isolated SCADA networks while conducting site surveys for DoD SCADA defense experiments. Network personnel had installed the wireless devices to facilitate remote monitoring. In both cases, one government and the other commercial, the wireless router installations were not reflected in the architecture diagrams and supervisory personnel were unaware of their existence. Wireless access, authorized or not, introduces a myriad of opportunities for human error and malicious access to what NIST and NERC-CIP specify should be an external connection free space. Even a logically isolated SCADA network is still vulnerable to attacks that jump the logical separation barrier, as was done with *StuxNet* and *Black Energy* attacks. Currently, cyber-defense of control systems relies heavily on the control system's human maintainers to observe system screens or to intuit deviations from normal operating parameters.

Expected DoD Benefit: More resilient and intrusion tolerant power generation at DoD installations during cyber-attacks. Improved continuity of power and services to critical Command and Control (C2) and mission support areas during crises. Additionally, the improvements in the ability to defend SCADA systems against Byzantine Intrusion attacks could be directly employed by the NAVFAC, a partner in this project, to protect the power grids supporting critical shore-based infrastructure and operations.

1.2.1 Phase 1: Demonstration at PNNL

The technologies demonstrated initially were two NIDS (a machine-learning sensor and CTAM-R) and two OT network defense technologies (Spines and Prime).

The machine-learning sensor (MLS) / NIDS was located inside the firewall of the OT network. The MLS also incorporated an anomaly detection system. Current ICS defense practices have no NIDS inside the OT networks and rely on operator observations to determine if the system is under attack. The benefit of placing a NIDS inside a vulnerable and currently undefended area is increased cyber situational awareness and an ability to “see” into what had been previously an essentially dark, although critically important, space.

This sensor was unique and groundbreaking both for its location inside the normally unprotected OT network and for its use of machine-learning and anomaly sensing to detect network misbehavior. This MLS evolved into the MANA™ in Phase 2.




Figure 1. The MANA™ Process

CTAM is a database designed to provide tailored threat intelligence information. It was paired with the Radiflow signature and anomaly NIDS to form the “CTAM-R”. CTAM-R is a prototype advanced, intelligent IDS for ICS that compiles threats, TTPs, and detectable indicators from various, open sources into a single, coherent, ICS-specific database of threats. CTAM correlates detected events with associated TTPs and threats to guide user response, enable deeper analysis, and make new associations. CTAM compiles the threats it has detected and in turn shares those confirmed detections externally using STIX and Trusted Automated Exchange of Indicator Information (TAXII) formats. The Radiflow IDS uses modified versions of Suricata and Bro IDS, combined with advanced, tailored, and modular smart packet inspection to match network traffic data to provided indicator rules. Radiflow also uses anomaly-based detection to identify variations from a derived baseline of normal network behavior, enabling the detection of some previously unknown attack patterns. The CTAM-R Correlates detected events with associated TTPs and threats to guide user response, enable deeper analysis, and make new associations.




Figure 2. CTAM Data Flow

Spines is a commercially proven communications overlay for the internet that guarantees traffic timeliness and completeness as long as a physical route exists between sender and receiver. This is Spines first use within an OT network and is the first known use of a communications overlay on an OT network. Current practices have nothing inside an OT network to guarantee that data is complete and delivered within time standards except the inherent capacity of the network's architecture. Spines is a unique and powerful application that meets data surety requirements within an OT network.

Prime is a prototype multi-compiler that protects key applications against direct attack through the use of multiple replicas of the defended application each compiled differently to present a diverse defense, and each communicating with the other replicas to maintain system state. The attacker has to defeat the multiple replicas simultaneously with separate attacks, a statistically difficult challenge to meet. This project was Prime's first use outside of a laboratory or within an OT network. Current practices have nothing inside an OT network to protect devices or applications. Prime is a unique and powerful application that prevents loss of key applications or devices within an OT network.

Symbiote (Red Balloon Securities) prevents malware and other cyber-attacks from hijacking, disrupting or corrupting any embedded device (VOIP phone, network printer, etc.). It defends devices without requiring changes to source code or hardware design and without impacting the functionality or performance of the device. This technology has considerable commercial potential because it is highly effective within any type of embedded device environment, from consumer electronics to factories, connected cars and even power plants. Because of Symbiote's commercial success, Red Balloon Security withdrew from the ESTCP project after award. However, we did have a networked printer defended by Symbiote in the test control system network at PNNL.

1.2.2 Phase 2: Demonstration at HECO Honolulu Plant

In the second phase of our project, external constraints, and brilliant insight by our Johns Hopkins University partners, left us with only "two" remaining technologies, MANA™ and Spire.

Spire is the merger of the successful JHU Spines and Prime from Phase 1 PNNL testing, integrated with a hardened open source SCADA Master, a Programmable Logic Controller (PLC) / Remote Terminal Unit (RTU) Proxy, pvbrowser and OpenPLC. There are no existing integrated systems like this in current use. All components of Spire are open source programs hardened by JHU to resist attacks and to be better able to support Prime replication and to integrate with Spines.

The Spire fault and intrusion tolerant technologies were developed under the assumption that the OT network they would be installed in was already compromised, and that critical control system components needed to resist or overcome attacks and to continue to function despite adversary efforts to disrupt them.

The MLS was upgraded from the PNNL demonstration based on SME input and results. It was given an improved graphical user interface, changes were made to the correlation methodologies for better detection of threats, and an improved user display with alerts grouped by confidence level for ease of plant engineer identification and prioritization was added. The improved product was named MANA™ (Machine-learning Assisted Network Analyzer). The single display enables the engineers, and later incident responders, to make quick decisions while offering a single console displaying all attack information. This single display of prioritized data is linked to the attack data on the MANA™ to enable security analysts to drill down and discover insights and patterns.

1.3 OBJECTIVE OF THE PROJECT

The objective of this project was to implement, and successfully demonstrate, an aware, fault and intrusion tolerant cyber defense-in-depth of a power plant's control system network. A power plant operating with the cyber defense-in-depth technologies used in this project will be cyber-secure, Byzantine fault tolerant, and aware. A plant thus protected will be able to withstand nation state equivalent cyber-attacks. Demonstrating this proof of concept onsite in a commercial utility company power plant, the project shows a practical cyber-secure path forward for DoD, DHS, and commercial utilities. This project is also a proof of concept to the DoD that in cyberspace, as in the other domains of war, a defense-in-depth of carefully selected and layered capabilities is more resilient when attacked and is intrusion tolerant in its abilities to seal off and eliminate penetrations.

NOTE: The technologies are usable on any control system. In selecting the technologies, consideration was given to the technologies being system and device agnostic and therefore compatible with most systems.

Objective of the Demonstrations

Two field demonstrations were conducted during this project. The first field demonstration was conducted in a lab at PNNL on an OT machine-in-the-loop emulation in March and April 2017. This PNNL demonstration was both an experiment and a risk reduction event and had multiple objectives:

- To determine if the technologies performed as their developers advertised.

- To determine if the technologies could be integrated and would work together without inducing problems in each other.
- To determine if the technologies could be integrated into the OT technology machine-in-the-loop emulation without inducing problems to OT systems or equipment. The emulation was designed by PNNL, using input from HECO, to be as close to the OT systems in a HECO power plant as possible, within budget and proprietary software constraints.

As a risk reduction event (RRE), this demonstration proved that it is safe for the technologies to be installed in an actual power generation plant without risk to plant operators, equipment, or systems and also without inducing latency or other issues to the control system network. This finding was key to HECO permitting us to bring the technologies into a power plant.

The second field demonstration was an operational validation of the technologies conducted inside a functional, but offline, HECO power plant. The objective of the operational validation was to implement and successfully run an aware, fault and intrusion tolerant cyber defense-in-depth in a commercial utility company's power plant's control system network without affecting the OT network, its devices or the field devices controlled by the OT network.

Issues Validated at the Demonstrations

The Phase 1 experiment at PNNL established that the fault and intrusion tolerant system (Prime running within a Spines overlay) enabled the OT network to continue to function despite being under attack by a sophisticated cyber threat. The Sandia National Laboratories cyber red team was unable to penetrate the secure Spines communications overlay protecting information exchanges between the SCADA Master and controlled endpoints during the designed experiment runs. Given root access to a Prime replica, Sandia was unable to compromise other systems or affect communications between replicas. The machine-learning NIDS was able to see 72% of all red team activity. CTAM only saw 41% of the attacks. There was not sufficient time for Sandia to attack or compromise the Symbiote defended printer, so it was not tested.

Only two technologies of the multiple technologies used in the PNNL risk reduction experiment made it to the operational validation in the HECO power plant – Spire and the MANATM NIDS. The Phase 2 operational validation reinforced the Phase 1 results that an aware cyber defense-in-depth using the Johns Hopkins' fault and intrusion tolerant technologies and an aligned machine-learning NIDS provides significantly better protection for control systems than the NIST compliant base-line architecture alone. HECO's plant network used the 18 control families of NIST SP 800-53. In order to test the cyber defense-in-depth technologies many of the regulatory access, physical, and training controls were set aside (e.g., Access Control, Awareness and Training, Audit and Accountability, Contingency, Identification and Authentication and Incident Response, Maintenance, Media Protection, System and Services Acquisition, and Program Management). As in Phase 1, compliance with the cyber specific portions of NIST offered no effective protection from a sophisticated threat.

The results of the Phase 2 operational validation (Live Demonstration) establish that fault and intrusion tolerant technologies can protect a utility's OT networks from an advanced and determined cyber threat with insider access. The operational validation also demonstrated that the cyber defense-in-depth technologies function in a power plant, are compatible with power plant equipment, and do not induce latency or errors in plant networks or devices.

The *Black Energy* and *StuxNet* incidents show that even isolated ICS networks are subject to penetration when threats jump across air-gaps into the OT networks. The results of the Phase 2 operational validation strongly suggest that the DoD, the DHS, and commercial utilities should consider implementation of cyber defense-in-depth architectures using fault and intrusion tolerant technologies combined with an aligned machine-learning NIDS to protect ICS and SCADA systems in critical infrastructure.

1.4 REGULATORY DRIVERS

While not explicitly requiring new technology, the following Presidential Policy Directives and Executive Orders identify gaps and needs that the critical energy infrastructure cyber defense-in-depth fills:

1. Presidential Policy Directive 21 - Critical Infrastructure Security and Resilience, 12 February 2013.
2. Presidential Policy Directive 41 - United States Cyber Incident Coordination, July 26, 2016.
3. Executive Order 13636 - Improving Critical Infrastructure Cybersecurity, 12 February 2013.
4. Executive Order 13691 - Promoting Private Sector Cybersecurity Information Sharing, 13 February 2015. [NOTE: CTAM was brought in to fill this need.]
5. Executive Order 13800 - Strengthening the Cybersecurity of Federal Networks and Critical Infrastructure, 11 May 2017.

Executive Order (EO) 13636, Improving Critical Infrastructure Cybersecurity, calls for the development of a voluntary Cybersecurity Framework that provides a “prioritized, flexible, repeatable, performance-based, and cost-effective approach” for assisting organizations responsible for critical infrastructure services to manage cybersecurity risk. The EO was the impetus for the NIST Cybersecurity Framework which establishes five core functions for organizing basic cybersecurity activities at their highest level: Identify, Protect, Detect, Respond, and Recover. The critical energy infrastructure cyber defense-in-depth meets all of these higher order functions of cybersecurity:

- The critical energy infrastructure cyber defense-in-depth Protects the control system network.
- The NIDS Detects threats entering the control system network. Another technology provides endpoint defense that Detects threats attempting to tamper with or change SCADA devices.
- These technologies send alerts that aid in the Identification of threat. The NIDS protects the entrance to the control system network and flags specific traffic so human network defenders can isolate and analyze exactly what was observed by the NIDS.
- The Prime system Responds to attacks by ignoring the attacks and then reclaiming compromised systems, by Recovering any compromised system to a previous known good status synched with the state of the systems that remained uncompromised and online.

The critical energy infrastructure cyber defense-in-depth also meets ISO/IEC 27001, Control A.10.8.3, which focuses on protection of data during transportation/transmission to achieve confidentiality, integrity and availability. This is a Subcategory of the “Data Security” Category in the “Protect” Function.

Page Intentionally Left Blank

2.0 TECHNOLOGY DESCRIPTION

2.1 TECHNOLOGY OVERVIEW

2.1.1 Operations Network Defense

2.1.1.1 Spire System

Spire is an intrusion-tolerant SCADA system built with the successful Spines and Prime technologies demonstrated in Phase 1 at PNNL, further integrated with a hardened open source SCADA Master, Programmable Logic Controller (PLC) / Remote Terminal Unit (RTU) Proxy, pvbrowser and OpenPLC. There are no existing integrated systems like this in current use. All components of Spire are open source programs hardened by JHU to resist attacks and to be better able to support Prime replication and to integrate with Spines.

The Spire fault and intrusion tolerant technologies were developed under the assumption that the OT network they would be installed in was already compromised, and that critical control system components needed to resist or overcome attacks and to continue to function despite adversary efforts to disrupt them. Spire is designed to withstand attacks and compromises at both the system level and the network level, while meeting the timeliness requirements of power grid monitoring and control systems (on the order of 100-200ms update latency). The Spire system was created by Professor Yair Amir, Trevor Aron, Amy Babay, and Thomas Tantillo of JHU. It is currently developed by the Distributed Systems and Networks Lab at JHU (www.dsn.jhu.edu/spire/) and supported by Spread Concepts LLC. Spire includes a SCADA Master and PLC/RTU Proxy designed from scratch to support intrusion tolerance, as well as pvbrowser-based HMI, the Prime intrusion-tolerant replication engine, and the Spines intrusion-tolerant network overlay. The SCADA Master is protected by replication using the Prime intrusion-tolerant replication engine. Communication between Spire components is protected using the Spines intrusion-tolerant network. The Spire PLC/RTU proxy interacts with any device that uses the Modbus or DNP3 communication protocols over IP. OpenPLC is used to emulate PLCs that can be monitored and controlled by the system. The Spire software components employ diversity and proactive recovery to increase the resilience of the system by preventing a single exploit from compromising all instantiations of a particular software component and by periodically removing potentially undetected intrusions from the system. The system is designed to scale to multiple field devices, sites, and operations centers, but was restricted to a small implementation for ESTCP.

In the normal flow of the system, there are two main types of events: HMI commands and PLC/RTU updates. When an HMI command is initiated (e.g., an engineer make a change), the command is sent to the control-center SCADA Master replicas over the external Spines network. The SCADA Master replicas pass the command to their Prime daemons, which disseminate it to the data-center Prime daemons and execute a Byzantine-fault-tolerant agreement protocol to agree on the command. When the Prime daemons have agreed on the command, they pass it back to their SCADA Masters. The SCADA Masters then execute a threshold signing procedure on the command (so that the PLC/RTU proxy can verify that a sufficient number of replicas agreed on the command by verifying a single signature on a single message). The control-center SCADA Masters then send the threshold-signed command to the PLC/RTU proxies. The proxies verify that the command has a valid threshold signature and then pass it on to the PLC(s) or RTU(s).

2.1.1.2 SCADA Master

The SCADA Master is the core of a typical SCADA system. It collects data from the PLCs and RTUs in the field and sends commands to control the state of the field devices. The SCADA Master of the Spire system is designed from scratch to be intrusion tolerant. The SCADA Master is replicated using the Prime intrusion-tolerant replication engine (described below) to ensure that its functions continue to be carried out correctly as long as no more than a threshold number of SCADA Master replicas are compromised simultaneously. In general, $3f+1$ replicas can tolerate f compromised replicas without affecting SCADA Master performance (http://www.cnds.jhu.edu/pub/papers/TDSC_Prime_2011_07.pdf). In the ESTCP experiment at PNNL, four SCADA Master replicas were deployed and successfully tolerated one of the four replicas being compromised and used maliciously. At HECO, a broader model that allowed proactive recovery on the fly was used to validate the principal of proactive recovery during active operations. In the broader model, $3f+2k+1$ replicas can tolerate f compromised replicas and k recovering replicas simultaneously while maintaining system state and protection. Using this model, six replicas were used at HECO. This permitted one replica to proactively recover while another is compromised, leaving four ($2f+k+1$) fully functioning replicas.

2.1.1.3 Prime

Prime is a Byzantine fault-tolerant replication engine that protects the SCADA Master and provides meaningful performance guarantees even after some of the replication servers have been compromised. Prime was created at Johns Hopkins University by Yair Amir, Jonathan Kirsch, John Lane, Marco Platania, Amy Babay, and Thomas Tantillo. Prime uses the MultiCompiler developed by the UC Irvine Secure Systems and Software Laboratory to diversify the code layout of Prime servers in order to increase the resiliency of the system. The MultiCompiler uses a 64-bit random number to generate different variants of an application. Prime guarantees Safety (consistency of the correct replicas) and Liveness (the eventual execution of each update) as long as no more than f out of $3f+2k+1$ replicas are compromised, no more than k replicas are unavailable (e.g., due to crashes, network partitions, or proactive recovery), and the network is sufficiently stable. Prime additionally provides a stronger performance guarantee, Bounded-Delay. Bounded-Delay limits the amount of performance degradation that can be caused by malicious servers by bounding latency between replicas to stay within system required / operator set limits. Prime forces any lead replica that remains in power to meet a threshold level of performance, where the threshold is a function of the message delays between the correct servers in the system, which cannot be arbitrarily increased by the malicious servers.

Prime ensures that each update or change to the system (e.g., software update, HMI command, or new PLC/RTU data) is agreed upon by a sufficient number (at least $2f+k+1$) of SCADA Master replicas before being applied. This ensures that if no more than f replicas are compromised at the same time, malicious attacks - such as malware hidden in an update, will not be effective. Prime guarantees that the state of non-compromised SCADA Masters remains consistent, as long as the threshold number of tolerated compromises is not exceeded. In addition to guaranteeing consistency, Prime guarantees performance under attack. Even in the presence of a compromise, the delay for an update to be processed is bounded in time by the system's requirements. This capability is crucial to meet the needs of latency-sensitive applications in power plants.

Prime uses diversity and proactive recovery to substantially increase the effort required for an attacker to compromise multiple replicas simultaneously. In this way, if an adversary attacks all the servers in parallel, the probability to defeat more than f servers is low. Proactive recovery allows Prime servers to periodically rejuvenate to clean potentially undetected intrusions from the system. Combined with diversity, a totally different version of the Prime server is generated after each rejuvenation. Each replica in the system consists of an application replica paired with a Prime daemon co-located on the same machine. Prime daemons provide an ordering service for the application replicas, with each Prime daemon delivering a totally ordered stream of updates to its application replica. Application-specific state transfer must be handled at the application level. When a Prime daemon detects that a state transfer is needed (i.e., because it has missed message, for example due to a crash, partition, or proactive recovery), it will notify the application, which can then decide on the correct action (different applications may vary on whether an application-level state transfer is needed).

By default, Prime is configured to make use of Spines, an overlay messaging framework developed at Johns Hopkins. Spines offers an intrusion-tolerant network service that protects communication between the Prime replicas. Spines can be deployed in both local-area and wide-area environments and includes tools for emulating wide-area topologies in local-area networks by placing bandwidth and latency constraints on the links between servers.




Figure 3. Abstraction of How Prime Defends

2.1.1.4 pvbrowser-based HMI

Spire uses a customized and hardened HMI derivative of the open-source pvbrowser. Pvbrowse is an application framework that provides a specialized browser for the client computer and an integrated development environment for the data acquisition programs that connect to the SCADA Master.

2.1.1.5 PLC/RTU Proxy

The PLC/RTU Proxy is a hardened device that relays Modbus or Distributed Network Protocol (DNP3) commands between the PLC or RTU and the SCADA Master. Before relaying a command, the proxy verifies that the command has been agreed upon by a sufficient number of uncompromised SCADA Master replicas. PLCs and RTUs are designed for rugged work conditions without regard for cyber protection.

The proxy protects these vulnerable devices from cyber threats while also using the Spines software to ensure delivery of traffic within tight SCADA timeliness requirements.

2.1.1.6 Spines

Spines provides an encrypted, timely, and highly reliable intrusion-tolerant messaging framework that overlays the physical network infrastructure. It connects the HMI, SCADA Master replicas, and PLC/RTU Proxy devices. Spines encrypts traffic to provide confidentiality, authenticates traffic to ensure integrity, and enforces resource allocation to ensure availability, giving it the ability to overcome denial of service attacks by compromised Spines nodes. Spines also uses redundant dissemination to prevent compromised nodes from disrupting routing.

2.1.1.7 Diversity and Proactive Recovery

The Spire system uses diversity and proactive recovery to substantially increase the effort an attacker must expend to compromise multiple system components simultaneously. This is particularly important to prevent an attacker from compromising more than the tolerated threshold of SCADA Master replicas. All software components of the Spire system employ compiler-based diversity to make each instance of a software component look different to an attacker, preventing a single exploit from compromising all instantiations of the same software component (e.g., all the SCADA Masters). Diversity is achieved by compiling the software using a customized MultiCompiler derived from the MultiCompiler created at UC Irvine (<https://github.com/secursystems/multicompiler>).

Table 2 presents replica requirements across the top and scenarios down the left. The table illustrates the tradeoffs the different configurations make between the algorithmic sophistication and numbers of servers and sites they require on the one hand, and the safety and availability guarantees they provide on the other hand. The table includes an analysis of each configuration's ability to support proactive recovery, withstand a single disconnected or down site, withstand a single intrusion, and any combination of these three conditions.

Table 2. Replica Requirements for Guaranteed Performance

	1	2	1-1	2-2	4	6	4-4	6-6	3+3 (f=1,k=1), x+y	2+2+2 (f=1,k=1)	2+2+2+2 (f=1,k=2)	4+4+4 (f=1,k=4)	2+2+2+2+2 (f=1,k=3)	3+3+2+2+2 (f=1,k=4)	3+3+3+3 (f=1,k=4)	6+6+6 (f=1,k=7)
All Correct	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green
Proactive Recovery (PR)	Yellow	Green	Yellow	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green
Disconnected/Downed Site	Red	Red	Orange	Orange	Red	Red	Orange	Orange	Red	Green	Green	Green	Green	Green	Green	Green
Disconnected/Downed Site + PR	Red	Red	Orange	Orange	Red	Red	Orange	Orange	Red	Yellow	Green	Green	Green	Green	Green	Green
Intrusion	Gray	Gray	Gray	Gray	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green	Green
Intrusion + PR	Gray	Gray	Gray	Gray	Yellow	Green	Yellow	Green	Green	Green	Green	Green	Green	Green	Green	Green
Disconnected/Downed Site + Intrusion	Gray	Gray	Gray	Gray	Red	Red	Orange	Orange	Red	Red	Green	Green	Green	Green	Green	Green
Disconnected/Downed Site + Intrusion + PR	Gray	Gray	Gray	Gray	Red	Red	Orange	Orange	Red	Red	Yellow	Yellow	Blue	Green	Green	Green

Legend:

- Green = a fully operational system with performance guarantees under attack. Specifically, the system is guaranteed to process any update within the bounded amount of time necessary to support SCADA systems for the power grid (100-200ms).
- Gray = the system is not guaranteed to remain safe: a single intrusion can compromise the system state.
- Red = the system will remain safe but will not provide any guarantee of state since the network or a site will need to be recovered or repaired before progress can be made.
- Orange = the system will remain safe, but will not provide any guarantee of progress until the cold-backup control center is activated by the operator. (Orange > red).
- Yellow = a fully operational system with the exception that its performance guarantee is not met when a correct replica is undergoing proactive recovery.
- Blue = a fully operational system with the exception that its performance guarantee is not met in a very narrow specific case, where one of the two control centers is down or disconnected, there is an intrusion in the other control center, and the remaining correct server in that control center is currently undergoing proactive recovery.

The SCADA Master replicas and the Prime replication engine employ proactive recovery to prevent an attacker from exceeding the threshold number of tolerated compromises over time. Each replica is periodically rejuvenated by being powered off and restored from a known clean state, removing any potentially undetected intrusion. By cleaning intrusions from the system, proactive recovery forces an attacker to exceed the threshold of tolerated compromises within a short vulnerability window (i.e., the time between a round of rejuvenations), rather than allowing the attacker extended time to compromise the system.

Proactive recovery and diversity are combined such that each time a replica is rejuvenated it is recompiled to generate a new variant that has a high likelihood of being significantly different than any previous instantiation of that replica or any of the other replicas in the system.

2.1.2 Cyber Situational Awareness

2.1.2.1 Machine-learning Assisted Network Analyzer (MANA™) - Resurgo, LLC

The MANA™ Network Intrusion Detection System (NIDS) incorporates both machine-learning and anomaly-based detection technologies. The MANA™ NIDS uses multiple different algorithms during the training process for each type of sensing method, and the best performing algorithms for that network are chosen for employment. The MANA™ system employs the best performing algorithms simultaneously and can be configured to employ one algorithm per instance or many algorithms of different modality per instance (i.e., one NIDS can perform the function of several heterogeneous NIDS). The MANA™ system correlates all alerts internally and displays for the human network defender a list of alerts with confidence levels. Confidence levels are the statistical “surety” that the alerts shown are definitely attacks or severe network anomalies. Additionally, the MANA™ NIDS is passive. It receives all traffic via a network tap and is compatible with all networks using IP based traffic to include variations such as MODBUS over TCP/IP.

2.1.2.2 Cyber Threat Activity Matrix with Radiflow (CTAM-R)

CTAM is an advanced, intelligent intrusion detection system for ICS that integrates the CTAM and an IDS to support broad use within DoD control systems and many segments of critical infrastructure. CTAM is a database-based system, providing up to date threat intelligence information tailored to ICS threats including signatures and defined network activity. CTAM compiles ICS threats, TTPs, and detectable indicators from multiple, open sources into a single, coherent, ICS-specific database of threats. CTAM correlates detected events with associated TTPs and threats to guide user response, enable deeper analysis, and make new associations. CTAM compiles the threats it has detected and in turn shares those confirmed detections externally using STIX and TAXII formats. CTAM is unique in its ability to gather information from external sources for immediate use and simultaneously share new information externally with user specified partners, such as the National Cybersecurity and Communications Integration Center (NCCIC). Very few NIDS of any type have this ability to take in new threat intelligence from multiple sources and provide new threat information externally without human interaction at every step.

CTAM was paired with the Radiflow signature and anomaly sensor to form the “CTAM-R” IDS, a single system capable of servicing a wide range of ICS. The Radiflow IDS uses modified versions of Suricata and Bro IDS, combined with advanced, tailored, and modular smart packet inspection to match network traffic data to provided indicator rules.




Figure 4. Approximation of CTAM Analysis of a Detected Attack

Radiflow also uses anomaly-based detection to identify variations from a derived baseline of normal network behavior, enabling the detection of some previously unknown attack patterns. CTAM-R does both signature-based and anomaly-based threat detection by monitoring the traffic between nodes of the ICS network. Using signature-based detection, known threats are detected through the use of threat definitions or signatures in the form of Snort rules. Signature-based detection is managed primarily by CTAM collecting external threat intelligence and managing the rules active within the Radiflow subsystem. For anomaly detection, potential threats are identified by comparing current network traffic to a baseline model of expected (normal) traffic. To do this, Radiflow creates a model of expected traffic by monitoring control system network traffic over a period of time. This model is used by the sensors to detect anomalous traffic that potentially represents a threat. CTAM maintains a copy of model and shows this to users as a map of source and destination nodes. As designed, CTAM would typically reside on an internet connected environment, such as an enterprise network, while the Radiflow subsystem would reside on the ICS network, typically in an out-of-band system, semi-isolated system except for a link to the CTAM. Figure 7 shows the physical placement of the CTAM-R components in this type of employment.

2.2 TECHNOLOGY DEVELOPMENT

2.2.1 Operations Network Defense Technology Development

2.2.1.1 Spire System

When the JHU Computer Science Department (CSD) accepted the offer to team with Resurgo, LLC, in 2015 only the Prime intrusion-tolerant replication engine and the Spines network overlay existed. Spines was in commercial operation supporting major international broadcast companies and Prime was undergoing developmental testing within the JHU CSD labs. The JHU team saw the benefit of combining their two technologies to defend OT networks. They knew that the core of an OT network is its SCADA Master and that the SCADA master was extremely vulnerable and the highest payoff target for attackers to inflict maximum damage to a plant with minimal effort. The SCADA Master collects data from the PLCs and RTUs in the field, relays the data to the human operator and sends back commands to control the state of the field devices. Spire's SCADA Master is designed from scratch to be intrusion tolerant and is further protected by replication using the Prime intrusion-tolerant replication engine (described below) to ensure that its functions continue to be carried out correctly as long as no more than a threshold number of SCADA Master replicas are compromised simultaneously. In general, $3f+1$ replicas can tolerate f compromised replicas without affecting SCADA Master performance as shown in Table 2 (above) (http://www.cnds.jhu.edu/pub/papers/TDSC_Prime_2011_07.pdf).

Prior to June 2017, the Spire system did not exist as a system. Coming out of our March 2017 risk reduction event at PNNL, the team from Johns Hopkins University saw a need to modify and integrate the technologies tested at PNNL to provide a better defense for OT networks. Spire was designed with intrusion tolerance as its core principle. Spire protects the network and the SCADA Master while it provides an intrusion-tolerant consistent state.

Spire combines the proven components of SPINES and Prime with new open-source system components built from scratch and hardened to provide a complete top-to-bottom solution (<http://dsn.jhu.edu/spire>):

- Spines (<http://spines.org>) - Intrusion-Tolerant Network
- Prime (<http://dsn.jhu.edu/prime>) - intrusion-tolerant, Byzantine fault tolerant, replication that guarantees performance under attack
- Pvbrowser-based HMI (<https://pvbrowser.de/pvbrowser/index.php>)
- SCADA Master (<http://dsn.jhu.edu/spire>)
- PLC/RTU Proxy (<http://dsn.jhu.edu/spire>)
- OpenPLC (<http://www.openplcproject.com>)
- Multicompiler (<https://github.com/secursystems/multicompiler>)
 - Diversity (Michael Franz group at UC Irvine, Immunant)




Figure 5. Spire Architecture with a Single Control Center

The Spire software components employ diversity and proactive recovery to increase the resilience of the system by preventing a single exploit from compromising all instantiations of a particular software component and by periodically removing potentially undetected intrusions from the system. The system is designed to scale up to support multiple field devices, sites, and operations centers, but was restricted to a single site small implementation for ESTCP.

In the ESTCP experiment at PNNL, four SCADA Master replicas were deployed and successfully tolerated one of the four replicas being compromised and used maliciously. At HECO, a broader model that allowed proactive recovery on the fly was used to validate the principal of proactive recovery during active operations.

In the broader model, $3f+2k+1$ replicas can tolerate f compromised replicas and k recovering replicas simultaneously while maintaining system state and protection. Using this model, six replicas were used at HECO. This permitted one replica to proactively recover while another is compromised, leaving four ($2f+k+1$) fully functioning replicas.

2.2.1.2 Spines

Spines is also a product of the JHU Computer Science Department and provides encrypted, timely, and highly reliable, intrusion-tolerant messaging framework that overlays the physical network infrastructure. A Spines variant has been deployed commercially since September 2014 over the global cloud. Originally deployed as an experimental, it is now in full production use by multiple television networks. Spines connects the HMI, SCADA Master replicas, and PLC/RTU Proxy devices. Spines encrypts traffic to provide confidentiality, authenticates traffic to ensure integrity, and enforces resource allocation to ensure availability, giving it the ability to overcome denial of service attacks by compromised Spines nodes. Spines also uses redundant dissemination (sending data over multiple paths simultaneously) to prevent compromised nodes from disrupting routing.

Communication in the system occurs over two Spines overlay networks: an external network and an internal network. The external network is used for communication between the SCADA Master replicas and the PLC/RTU proxies and the HMIs. The internal network is used for communication among the SCADA Master replicas (and their Prime daemons). External and internal SPINES daemons can be deployed on the same machines but use different ports.

2.2.1.3 Prime

Prime is a Low Level Virtual Machine (LLVM) that runs an intrusion-tolerant replication engine that protects the SCADA Master. Prime began as a research and development project and has been in lab use since 2012. Integration of Prime with a JHU selected open source SCADA HMI was DARPA funded and completed in November 2015 at JHU. Prime ensures that each update to the system (e.g., HMI command or new PLC/RTU data) is agreed upon by a sufficient number (at least $2f+k+1$) of SCADA Master replicas before being applied. This ensures that if no more than f replicas are compromised at the same time, malicious attacks - such as malware hidden in an update, will not be effective and all normal functions will continue to occur within the operator specified bounded delay. NOTE: many functions on an OT network are time sensitive. Bounded delay is the specified maximum permitted latency that ensures complete functionality and timeliness. Prime guarantees that the state of non-compromised SCADA Masters remains consistent, as long as the threshold number of tolerated compromises is not exceeded. In addition to guaranteeing consistency, Prime is Byzantine fault tolerant and guarantees performance under attack. Even in the presence of a compromise, the delay for an update to be processed is bounded in time by the system's requirements. This capability is crucial to meet the needs of latency-sensitive applications in power plants.

2.2.1.4 pvbrowser-based HMI

Spire uses a customized and hardened HMI derivative of the open-source pvbrowser. Pvbrowser is an application framework that provides a specialized browser for the client computer and an integrated development environment for the data acquisition programs that connect to the SCADA Master.

Process Visualization Browser (pvbrowser)(pvh®) is an application framework created by Rainer Lehigh. It is an open source SCADA software solution that is platform-independent and runs on Linux/Unix/Windows/OS-X and OpenVMS operating systems. The pvbrowser software permits the user to choose the operating system which promises the best security. This is important where pvbrowser clients run in less secure environments, such as on OT networks.

The pvbrowser client is a C++ application on top of the Linux Qt graphical toolkit for the browser and ANSI C for the server, optimized for high update rates to support process visualization. The pvbrowser uses the connection oriented pv (Process Visualization) protocol. However, the normal connectionless http (Hyper Text Transport Protocol) protocol may also be used from pvbrowser. Thus, the user of pvbrowser can browse the distributed process visualization servers (pvservers) with an optimized protocol and can also view data served over the http protocol. The pvbrowser function permits technical processes to be visualized or controlled. All configuration is done on the server-side. The engineer / operator can modify the server, writing routines that interact with hardware, or software, and defines which objects are to be displayed and controlled in which way. These components are displayed by the pvbrowser according to the information provided by the pvserver over the OT network.

The pvbrowser software consists of a server (pvserver) and a browser (pvbrowser), which communicates with the engineer / operator via an HMI. The pvbrowser client can also be configured to talk to the server only via SSH or VPN if further security is required. The pvbrowser provides a structure for process visualization. It was selected because of its use in a Romanian power distribution system covering 10,000 km².

NOTE: The Process View Browser is only Free Software under GNU/Linux, for which it is released under the GNU General Public License. The pvbrowser is proprietary under Windows and VMS.

The pvbrowser software also contains an open source SCADA Master and HMI server.

Architecture

- SCADA Master
 - Data acquisition daemon (DAD): communicates with RTUs/PLCs
 - Shared memory: medium for communication between DAD and pvserver
 - ProcessViewServer: visualizes data from DAD and communicates with HMI
- HMI
 - ProcessViewBrowser (pvbrowser) presents information from the visualizer to use

PvBrowser was modified by JHU to make it more secure and to permit replication. Memory is no longer shared between the daemons (DAD) and the pvserver, and messaging was changed to permit direct message passing between DAD and the pvserver. This eliminated nondeterministic behavior and allowed processes to run on different machines. It also enables replication to be used.

DAD and 3f+1 PVServers all run on different physical or virtual machines

- DAD polls RTUs and communicates data to servers via message passing
- Data polled from field units is replicated consistently across all servers using Prime
- The HMI can connect to any server and observe consistent data at each correct replica

This provides the intrusion tolerant replication of the data used by the SCADA Master.

2.2.1.5 SCADA Master

The SCADA Master was custom built by JHU in late 2017 using the framework provided by the SCADA Master in pvbrowser in order to have a more secure software that performs better with Prime. Each actual SCADA Master is replicated by Prime. For ESTCP all replicas were at the same physical location but for optimal resiliency they would be distributed over multiple physical sites connected by a wide-area network, over multiple logical sites within a local-area network (with or without emulated wide-area latency) or as a single site in a local-area network. Ideally, Spire would be deployed with its SCADA Master replicas distributed across several sites.

For each SCADA master replica, a Prime daemon is deployed that the SCADA master connects to. Each SCADA master is located on the same physical machine as its Prime daemon and connects to the Prime daemon via interprocess communications. Only two types of sites can contain SCADA Master replicas: control centers and data centers. This is because power grid control centers with full capabilities for controlling PLCs and RTUs are expensive, and utility companies are unlikely to deploy more than two. To support the desired resilience with only two control centers, additional sites can be added as data-center sites that do not control PLCs or RTUs. In each site that contains SCADA Master replicas (including both control centers and data centers), one Spines daemon is deployed that participates in the internal network to connect the replicas in that site to the other sites. In each control-center site, an additional Spines daemon is deployed that participates in the external network to connect the replicas in that site to the proxies and HMIs.

2.2.1.6 PLC/RTU Proxy

A PLC is an industrial digital computer which has been ruggedized and adapted for the control of manufacturing processes, such as assembly lines, or robotic devices, or any activity that requires high reliability control and ease of programming and process fault diagnosis. A PLC is a "hard" real-time system since output results must be produced in response to input conditions within a limited time, otherwise unintended operation will result. The PLC/RTU Proxy is a JHU hardened device that secures the PLC/RTU and relays Modbus or Distributed Network Protocol (DNP3) commands between the PLC or RTU and the SCADA Master. Before relaying a command, the proxy verifies that the command has been agreed upon by a sufficient number of uncompromised SCADA Master replicas. PLCs and RTUs are designed for rugged work conditions without regard for cyber protection. The proxy protects these vulnerable devices from cyber threats while also using the Spines software to ensure delivery of traffic within tight SCADA timeliness requirements. PLCs and RTUs are periodically polled by their proxies. When a proxy has new PLC/RTU data, it sends the data to the control-center SCADA Master replicas over the external Spines network to be agreed upon and sent to the HMI. The HMI verifies the threshold signature on the update and updates its display.

2.2.1.7 Diversity and Proactive Recovery

The Spire system uses diversity and proactive recovery to substantially increase the effort an attacker must expend to compromise multiple system components simultaneously. This is particularly important to prevent an attacker from compromising more than the tolerated threshold of SCADA Master replicas.

All software components of the Spire system employ compiler-based diversity to make each instance of a software component look different to an attacker, preventing a single exploit from compromising all instantiations of the same software component (e.g., all the SCADA Masters). Diversity is achieved by compiling the software using a customized MultiCompiler derived from the MultiCompiler created at UC Irvine (<https://github.com/securesystemslab/multicompiler>).

The SCADA Master replicas and the Prime replication engine employ proactive recovery to prevent an attacker from exceeding the threshold number of tolerated compromises over time.

Each replica is periodically rejuvenated by being powered off and restored from a known clean state, removing any potentially undetected intrusion. By cleaning intrusions from the system, proactive recovery forces an attacker to exceed the threshold of tolerated compromises within a short vulnerability window (i.e., the time between a round of rejuvenations), rather than allowing the attacker extended time to compromise the system. Proactive recovery and diversity are combined such that each time a replica is rejuvenated it is recompiled to generate a new variant that has a high likelihood of being significantly different than any previous instantiation of that replica or any of the other replicas in the system.

2.2.2 Cyber Situational Awareness Technology Development

2.2.2.1 Machine-learning Assisted Network Analyzer TM - Resurgo, LLC

The MANATM NIDS began as an unnamed side project by Resurgo programmers. In our original proposal, Oak Ridge National Laboratories had agreed to partner with us and provide their machine-learning sensor. On award, Oak Ridge significantly increased how much money they wanted to participate in the cyber aware defense-in-depth. After contacting other companies with machine-learners, Resurgo decided to self-fund the development of a machine-learning NIDS to meet the obligations of our proposal. Resurgo's development of MANATM was influenced by a DARPA project "Automating the Training of Machine-learning Sensors" for which Resurgo served as prime contractor. Machine learning is the eyes and brains of the defense-in-depth. In a six-month period, using open source components, Resurgo's all cleared U.S. citizen programmers developed, modified and applied for a patent for MANATM. MANA's first use was at PNNL as the MLS. Based on the MLS's performance, and suggestions from the SMEs at the event, a human readable graphical interface (dashboard) was developed that combined all automated data, scrubbed and correlated the data, and displayed it in a manner that facilitated cognitive reasoning and quick response by plant engineers without a cyber background. Alerts are displayed by confidence level, the probability that a given alert is in fact an attack or exploit, with high medium and low confidence categories to provide context. This enables network defenders or incident responders to make proper decisions while offering them an explicit visual console. The MANATM interface is a single display of relevant data that enables a security analyst to drill down and discover insights and patterns with access to 100% of the network traffic and machine alerts that went into the production of the displayed correlated alerts.




Figure 6. MANATM Employment Schematic

2.2.2.2 Cyber Threat Activity Matrix with Radiflow

The initial prototype of CTAM was created in response to solicitation W912BU-14-R-0023 a requirement under the NORTHCOM-13-BAA-RIF-0001. This led to a contract for an Industrial Control Systems Cyber Intrusion Detection System, with a period of performance from 24 Sep 2014 to 15 Jun 2016. The initial prototype was developed by Imprimis, Inc., partnered with subcontractor root9B.

Imprimis focused on developing the CTAM integrated cyber threat database, post-alert analysis, user interface, and overall project management. CTAM compiles threats, TTPs, and detectable indicators from various, open sources into a single, coherent, ICS-specific database of threats. The CTAM correlates detected events and associates them with TTPs and threat information in its database to guide user response, enable deeper analysis, and make new associations. It was designed to support a NIDS, supplying the NIDS with the latest signatures and threat information to both identify the attack and the attacker. CTAM compiles the threats it has detected and in turn shares those confirmed detections externally using STIX and TAXII formats to the corporate headquarters of the user, to partner companies and to the DHS's NCCIC.

Root9B developed the Smart Packet Inspection IDS (SPI/IDS) subsystem, which combined an advanced, tailored, and modular smart packet inspection (SPI) anomaly detection tool combined with modified versions of the open source IDS, Suricata and Bro. CTAM SPI/IDS used a master/slave architecture, where CTAM facilitated external oversight and control over the SPI/IDS ICS monitoring subsystem. SPI/IDS monitors the ICS network and issues alerts according to a cyber monitoring plan (CMP) received from CTAM. The NORTHCOM customer arranged for the CTAM-SPI/IDS system to be fielded on a cyber range at the US Air Force Academy (USAFA).




Figure 7. CTAM-R Employment Schematic

In negotiating the contract to support the ESTCP PNNL event, the initial intent was to use CTAM SPI/IDS successfully proven at USAFA. However, funding and timeline limitations dictated a consideration of other options for CTAM employment and a decision to demonstrate CTAM flexibility and extensibility by using a different sensor subsystem. After analyzing several options, Imprimis partnered with Radiflow based on their expertise with industrial control systems. CTAM would integrate with Radiflow sensors and their industrial IDS (ISID) as a sensor controller. CTAM still provided cyber threat information, a user interface and overall process control, as well as providing the system log event data to the provided Security Information and Event Management (SIEM) like capability gained by integrating Elastic Search, Logstash, and Kibana (ELK) applications. Expanding CTAM to integrate with Radiflow components, and provide a specific SYSLOG feed to the SIEM required several primary changes to CTAM, primarily related to differences in how SPI/IDS and Radiflow ISID communicate, and in how they treat network maps and alerts. The signature component of the Radiflow ISID uses modified versions of Suricata and Bro IDS, combined with advanced, tailored, and modular smart packet inspection to match network traffic data to provided indicator rules. The Radiflow ISID also uses anomaly-based detection to identify variations from a derived baseline of normal network behavior, enabling the detection of some previously unknown attack patterns. The CTAM-R Correlates detected events with associated TTPs and threats to guide user response, enable deeper analysis, and make new associations. The original CTAM-SPI/IDS capabilities all remained in place, with newly created code used to execute Radiflow-specific capabilities. The primary differences included:

- **Alternative Message Handling.** The original CTAM SPI/IDS version made extensive use of the TAXII protocol as the primary communication mechanism between CTAM and SPI/IDS. This system worked well and helped ensure standardized formatting. However, it did require the use of multiple TAXII servers and a fair amount of processing, time, and overhead. Additionally, it would have necessitated considerable changes to ISID to get it to communicate via TAXII. Therefore, in CTAM-R, CTAM and ISID instead communicate via standardized Representational State Transfer (RESTful) application program interface (APIs). This provided considerable flexibility, minimized necessary software changes, and proved more efficient. Additionally, this lays the groundwork for CTAM to be extended to work with other systems that communicate via APIs
- **Alternative Baseline Handling.** SPI/IDS and ISID both create a baseline of expected network behavior by observing activity over a period of time – a process of learning or baselining. In general, both model observed nodes, links between nodes, and characteristics of communication between nodes. However, the details of these models are quite different, focusing on different aspects structuring the model in very different ways. Additionally, the way the systems operate when in learning or baseline mode are quite different. Therefore, in CTAM-R additional interfaces and options related to managing system learning had to be added, as well as different mechanisms to allow the user to accept certain anomalous behavior as ‘normal’ and expand the model.
- **Alternative User Displays.** SPI/IDS and ISID provide different details in the baseline model and in detected events. This necessitated adding alternative user displays, to show the ISID-provided data vice SPI/IDS-provided data, or adding options to specific user displays, based on whether CTAM was connecting to SPI/IDS or ISID. Most of the user displays are the same, but where they are different, if CTAM connects to SPI/IDS, it will use SPI/IDS-specific displays, if ISID, then ISID-specific displays.
- **Enhanced Logging Capability.** Finally, interfacing with Resurgo’s ELK SIEM required the extension of CTAM logging capabilities. In CTAM SPI/IDS, CTAM created numerous internal logs stored in the CTAM database. However, they were all for internal CTAM usage, or user review, not for export. Therefore, in CTAM-R the team added capability to write to external SYSLOG servers in a specific format, initially based on a standard SYSLOG format, but then refined in coordination with Resurgo to support their event data-gathering requirements.
- **Handling Snort IDS Rules.** SPI/IDS utilized modified versions of the Suricata and Bro open source IDS’ and made extensive use of Suricata detection rules for signature detection. ISID instead utilizes a modified version of the Snort IDS engine and, therefore, requires the use of Snort detection rules. This necessitated the extension of CTAM cyber threat source importer functionality to import Snort rules as well as modifications to the processes to allow for rule synchronization between CTAM and Radiflow.

2.3 ADVANTAGES AND LIMITATIONS OF THE TECHNOLOGY

There are no existing or comparable technologies for the fault and intrusion tolerant defense of control system networks or SCADA devices. There are, however, other machine learning sensors.

The majority of these other machine learning sensors are designed for traditional IT networks. Those that have been modified to run on OT networks tend to be single algorithm NIDS without correlation, unlike the MANA™ NIDS that uses a proprietary mix of algorithms and uses proprietary correlation methodology. There are also comparable technologies; the signature and anomaly NIDS used on most business IT networks. Anomaly NIDS work by finding things outside of administrator defined “normal” parameters or observed “normal” traffic. For example, if office systems are used from 8 AM to 6 PM, Monday through Friday, an anomaly detector would detect and report use at 10 PM on Wednesday or any time Saturday or Sunday. An authorized user, working at an unusual time on authorized work, will set off an anomaly NIDS. Signature sensors, like McAfee or Symantec, rely on the software pattern of previously successful attacks as a pattern that is matched against all current traffic. If the previous attack is altered in any way, the signature sensor will not find a pattern match and will not be able to detect it. An MLS is trained using carefully screened “good” or “normal” traffic for a network. The MLS is then trained on the types of attacks typically seen on the same network. This gives the MLS a basis to judge good and bad similar to teaching a young child what is permitted and what is not. When the MLS is exposed to something it has never seen before in traffic, it uses its training to evaluate the traffic, and the traffic’s behavior, and it judges the traffic. It will flag the traffic and report the traffic as bad for a network defender to evaluate if it has any “doubts” about the traffic.

- Performance Advantages:
 - The Spire technologies were judged by HECO engineers to decrease latency on the portion of the network Spire protected by over 40%.
 - The MANA™ NIDS using a proprietary mix of algorithms and proprietary correlation methodology to deliver high confidence alerts on a graphical user interface (GUI) in human readable format.
 - The MANA™ NIDS is positioned out of band and is passive, creating no latency or burden on the network it is defending.
- Performance Limitations: Some utility network components, such as PLCs, must perform within tight specifications that require high reliability control, fault diagnosis, and sometimes ease of programming. The PLC Proxy functioned in near real time with no discernable latency at PNNL or at HECO as judged by HECO engineers independently.
- Cost Advantages: There are significant cost advantages to the current cyber best practices
 - they expend few labor hours or dollars. The DoD and utility companies pay only for minimal recommended software and equipment for their defenses and benefit from cost avoidance in never having suffered an equipment damaging attack or being held responsible for outages. However, based on Russian cyber activities against Ukrainian power plants, this will change and the costs to the power provider in lost business and damaged equipment will quickly swing the cost advantage to ANY defense, regardless of the acquisition cost. The cyber defense-in-depth provides:
 - Cost savings as a byproduct of increased availability of services vs dollars and lives lost due to mission failure if power generation and services are disrupted.

- Additional cost avoidance from not encountering potentially large replacement costs for turbines and other power generation infrastructure components physically damaged by a cyber-attack.
- Cost Limitations: Initial acquisition costs are approximately \$250,000 for the entire suite of technologies, including installation and travel to the installation site. This cost estimate is to obtain all of the technologies in low production at no discount. Initial acquisition costs at low production rates are higher than current signature or anomaly NIDS. Annual defense-in-depth technology costs are projected to be approximately the same as a subscription with support contract for the current, less effective, NIDS.

Operations and maintenance costs are site specific depending on the salary costs of the installation site's network defenders. None of the technologies will require the hiring of additional workers to support operations after they are installed and operating. The MLS does require extensive training and purchasers are advised to get a support contract to update (retrain for newer threats) the MLS as required. The Spire system while offered as open source software has a more powerful version with a support contract for purchase.

The alternative technologies currently in place are the traditional signature and anomaly based intrusion detection/intrusion protection systems used on the Enterprise/business information technology networks. In our Phase 1, PNNL experiment, Sandia National lab's cyber red team actions went unobserved by the signature based sensor which only caught one of 53 attacks. Because every power plant is unique, running older equipment and using proprietary systems, Sandia had to craft individual attacks that had no known signatures. Sandia's attacks also hid in the machine traffic and were almost unseen by the anomaly based sensors, one of which caught 66% of Sandia's attacks and the other caught less than 30%.

There are multiple cost benefits to the current practices and technologies; they make use of the same systems and skills for the OT network that the IT network uses, they are minimal in additional cost and they have worked to date but only because they have not been subjected to the sophisticated threat levels manifested in this ESTCP effort.

3.0 PERFORMANCE OBJECTIVES

3.1 PHASE 1 - PNNL RISK REDUCTION (LABORATORY)

3.1.1 Phase 1 - PNNL Risk Reduction (Laboratory) Performance Objectives

Table 3. PNNL Risk Reduction Performance Objectives

Performance Objective	Metric	Data Requirements	Success Criteria
Quantitative Performance Objectives			
Architecture Cyber Security	Data Stream Protection Score	Confidentiality Score	Contents unknown to Red Team
		Integrity Score	Contents unchanged by Red Team
		Availability Score	Contents unhindered by Red Team
Sensor Performance	Sensor Precision	True-Positive Alert	Correctly identified Red Team activity
		False-Positive Alert	Categorized normal traffic as Red Team activity
	Sensor Recall	Caught Attack	Correctly identified Red Team attack
		Missed Attack	Non-identified Red Team attack
Qualitative Performance Objectives			
Control System or Plant equipment damage	Cost of repairs or replacement (US Dollars (\$))	Test information assessed by SMEs from HECO, PNNL, and NAVFAC *	No damage to systems or induced issues as determined by SMEs.

*The Subject Matter Experts (SME) used their knowledge of the systems, attacks, and access to attack truth data to verify CEH claims and to assess probable damage to systems(s) or equipment.

3.1.2 Phase 1 - PNNL Risk Reduction (Laboratory) Performance Objectives Descriptions

A. Performance Objective 1 – Architecture Cyber Security:

- **Name and Definition:** Ability of the architecture's network defense systems to thwart cyber-attacks and to protect functionality of essential supervisory and control links.
- **Purpose:** To compare and contrast the current cyber defense standard for utility control systems to a potential, more complex iteration (i.e., defense-in-depth).
- **Metric:** Architecture cyber security was assessed by the aggregate score of all designated data stream scores, where K_n is a weighting factor that can be applied to data exchanges that are deemed more valuable by SMEs.

$$\text{Aggregate Score} = \sum_{n=0}^{\text{\# of Data Exchanges}} [K_n \times \text{Data Exchange Score}_n]$$

For individual data exchanges, we used a scoring system based on the confidentiality, integrity and availability with assigned values (high, medium, low) given to the supervisory control and data acquisition (i.e., monitor) components.

$$Data\ Exchange\ Score = Score_{Confidentiality} + Score_{Integrity} + Score_{Availability}$$

$$Score_{Confidentiality} = Score_{Monitor} + Score_{Control}$$

$$Monitor\ Score = \begin{cases} 3\ points\ | \ high \\ 2\ points\ | \ medium \\ 1\ point\ | \ low \end{cases}$$

The “white cell” gave a score to each data exchange relative to how these requirements were fulfilled during test runs. Each individual data exchange score was aggregated to create a value representative of the whole system (i.e., cyber-secure utility architecture):

Example: A simple system, which only has a few exchanges, could have the below assigned scores for communication between an HMI and an RTU, where CIA stands for Confidentiality, Integrity and Availability:

Table 4. CIA Scoring Example

Information Assurance Requirements	Monitor	Control
Confidentiality	Low(1)	Medium(2)
Integrity	Medium(2)	High(3)
Availability	Medium(2)	High(3)
Maximum Score	5	8

There are two responsibilities for the SCADA data exchanges: monitor and control. The minimum score for the architecture protecting this exchange is 0 (if the red team violates all confidentiality, integrity and availability requirements) and the maximum score is 13 (if the red affects nothing). If the SMEs decide that this communication is twice as important as other data exchanges, a weight factor ($K_n = 2$) could be applied to the calculation, for a maximum possible score of 26 for this data exchange. Aggregate score for the architectures’ cyber security would be 26 (for this exchange) plus the other data exchange scores.

- Data: A “white team” with full access to red team activity and all network data information grade the architecture’s CIA protection of each data stream after a test run.
- Analytical Methodology: Architecture results were analyzed for statistically significant differences using the Paired T-Test or a similar parametric type test. When statistical analysis was not useful, a straight comparison of the control group (standard utility architecture) to the test group (defense-in-depth) was done.

- Success Criteria:
 - a. Confidentiality data exchange score of greater than 95% and contents unknown to red team.
 - b. Integrity data exchange score of greater than 95% and contents unchanged by red team.
 - c. Availability data exchange score of greater than 95% and contents unhindered by red team.

B. Performance Objective 2 – NIDS Performance:

- Name and Definition: NIDS performance provided a quantitative assessment of usefulness by looking at the percentage of correct alerts (i.e., precision) and the percentage of attacks actually caught (i.e., recall).
- Purpose: To evaluate the usefulness of various types of NIDS at various locations within the utility control system network to determine optimum location.
- Metric: By analyzing the NIDS event/alert reports after each test run, the white cell was able to categorize each alert as a true or false-positive based on the Internet Protocol (IP) addresses assigned to the red team. Similarly, for red team specific efforts, we were able to identify which attacks were caught or missed. These become our measures of performance and the basis for our measures of evaluation.

Alert Logs (Events)

- True-Positive Alerts (correct events)
- False-Positive Alerts (wrong events)

Attacks (Exploits)

- Caught Attacks
- Missed Attacks

NIDS recall and precision were calculated for each individual NIDS and after combining both NIDS results at the SIEM-level. The difference in results demonstrated the effects of implementing the heterogeneous NIDS.

$$Recall = \frac{Caught\ Attacks}{Missed + Caught\ Attacks} = \frac{Caught\ Attacks}{Total\ Attacks}$$

$$Precision = 1 - \frac{FalsePositives}{TruePositives + FalsePositives} = 1 - \frac{FalsePositives}{TotalPositives}$$

- Data: To score the NIDS alerts, the “white cell” needed to know the red team IP addresses.
- Analytical Methodology: For comparing NIDS to each other or to the combined NIDS system at the SIEM, we determined if the differences in results were statically

significant with an Independent T-Test. No analytical methodology was used for comparison with the control group (standard utility architecture) since intrusion detection is not a standard practice within control system networks.

- Success Criteria:
 - a. NIDS Precision:
 - i. True-positives = greater than 90% of red team attacks correctly identified.
 - ii. False-positives = less than 10% of normal traffic identified as red team activity
 - b. NIDS Recall:
 - i. Caught attacks = greater than 90% of red team attacks correctly identified.
 - ii. Missed attacks = attacks that did not identify red team activity, less than 10%.

3.2 PHASE 2 - HECO OPERATIONAL VALIDATION (Power Plant)

Cyber Defense-in-Depth Technology Performance Objectives:

Our performance objectives measured the impact of the technologies on plant operations and tolerances within the utility control system. Using HECO equipment, we determined if the technologies caused any latencies or changed plant output. (NOTE: since the plant's power was not required by HECO during testing, all igniters and boilers were offline.)

A certified ethical hacker (CEH) team provided the attacks and erroneous stimuli to the network. The CEH, under HECO guidance, focused on devices and protocols HECO is concerned about within the control system. The CEH ignored all extraneous networked platforms (e.g., printers).

3.2.1 Phase 2 - HECO Operational Validation (Power Plant) Performance Objectives

Table 5. Performance Objectives

Performance Objective	Metric	Data Requirements	Success Criteria
Quantitative Performance Objectives			
Compatibility	Any plant component’s failure	Position / condition of PLCs and RTUs	Continued component operation
	Any plant component’s performance degradation	Position / condition of PLCs and RTUs	Continued component operation within specified tolerances
NIDS Performance	Alerts	Caught Attack	Correctly identified CEH attack
		Missed Attack	CEH attack not identified
Qualitative Performance Objectives			
Control System or Plant equipment damage	Cost of repairs or replacement (US Dollars (\$))	Test information assessed by SMEs from HECO and NAVFAC *	Damage or induced latency as determined by the SMEs.

*The SME's used their knowledge of the systems, attacks, and access to attack truth data to verify CEH claims and to assess probable damage to systems(s) or equipment.

3.2.2 Performance Objectives Descriptions

Performance Objective 1 – Technology Compatibility:

- Name and Definition: Compatibility is the ability of all of the plant's systems and devices to continue to function normally.
- Purpose: To determine if the defense-in-depth technologies cause problems with or degrade the performance of any device or system within the plant's utility control systems.
- Metric: Observed changes in device performance or outputs from the baseline established PRIOR to any of the cyber defense-in-depth technologies being installed.
- Data: Observed or measured changes in any part of the power plant's performance.
- Analytical Methodology: Compare plant systems performance with the defense-in-depth technologies to the undefended portion of the plant's utility control system.
- Success Criteria:
 - a. The technologies do NOT induce latency or performance degradation.
 - b. The technologies do NOT induce equipment or device failure.
 - c. All plant systems continue to operate as they were prior to technology installation or operation.

Performance Objective 2 – NIDS Performance:

- Name and Definition: Sensor performance will be a pass / fail assessment of the sensor's ability to identify threats in a power plant.
- Purpose: To validate that the sensor performs correctly outside of the lab.
- Metric: Alerts recorded by MANA™.
- Data: To validate sensor alerts, the will need the CEH to provide:
 - The IP address from which the attack was launched
 - The IP address attacked
 - The time of the attack
 - The attack type / method
 - If the CEH received any feedback from the system after executing the attack
- Analytical Methodology: We will compare each event/alert reports after each test run to the red team's attack log and determine if the attacks were caught or missed.
- Success Criteria:
 - a. Caught attacks correctly identify greater than 90% of CEH attacks.
 - b. Missed attacks. NIDS does not fail to identify greater than 10% of the attacks.

Page Intentionally Left Blank

4.0 FACILITY/SITE DESCRIPTION

During the course of our project we used two facilities, PNNL and HECO, and performed two separate demonstrations.

4.1 FACILITY/SITE LOCATION AND OPERATIONS

4.1.1 Pacific Northwest National Laboratory

The Phase 1 demonstration was hosted by PNNL in Richland, Washington. PNNL provided an enclosed, unclassified industrial control system lab space within a secure building on its campus at 902 Battelle Blvd, Richland, WA 99354, where we worked undisturbed by other projects. PNNL is the preeminent Department of Energy (DoE) lab and was selected based on the recommendation of our sponsor, US Pacific Command (PACOM), with the concurrence of NAVFAC, our transition partner. HECO, the energy provider to PACOM, also concurred with the choice. PNNL replicated aspects of the HECO power plant either physically with the same components and software or by emulation of what it could physically replicate or obtain. The objective in designing the PNNL test architecture was to ensure fidelity in those portions of the network of particular interest to HECO.




Figure 8. The Pacific Northwest National Laboratory Campus

4.1.2 Hawaiian Electric Company Power Plant

The second demonstration, the live validation, was hosted by HECO in Honolulu, Hawaii. HECO is the island of Oahu's primary power provider and is the power provider for all DoD facilities on Oahu; including PACOM, one of our sponsors. The power plant used was offline awaiting decision to retire the plant or place it in reserve. All control systems and devices were functional but were last updated in 2017, so were vulnerable to newer cyber-attacks. All major components (boilers, pre-heaters, cooling systems, etc.) had been disconnected from the control systems for maintenance while awaiting decision on the plant's fate. This had the secondary effect of ensuring that no major systems could be inadvertently damaged during the demonstration.




Figure 9. The HECO Honolulu Power Plant

4.2 FACILITY/SITE CONDITIONS

4.2.1 PNNL

PNNL built a machine-in-the-loop emulation using specifications provided by HECO to the best of their ability given cost constraints. The site for the demonstration was inside a secure PNNL laboratory chosen for its ability to emulate energy producer critical infrastructure. All technologies and data gathering devices used in the demonstration were removed from the plant at the conclusion of the event.

4.2.2 HECO Power Plant

The operational validation was conducted in a licensed, functional, commercial electrical power plant. The portions of the power plant used for the operational validation were entirely enclosed. The plant was not air conditioned except for server rooms and control rooms for the plant operators. No additional environmental permits were required for the demonstration. At HECO's request, all hacking and attacks took place from physically within the power plant so that no external openings were required in their firewalls or existing cyber defenses. All technologies and data gathering devices used in the demonstration were removed from the plant at the conclusion of the event.

Page Intentionally Left Blank

5.0 TEST DESIGN

We performed two tests in this project and will first address the test at PNNL and then the operational validation at HECO.

5.1 PNNL CONCEPTUAL TEST DESIGN

Testing at PNNL involved a comparison of a standard utility's network architecture (emulated) to a layered defense-in-depth of the same emulated network architecture. The standard network architecture conformed to the Purdue Enterprise Reference Architecture (PERA) model and followed NIST recommended cyber security implementations. A cyber red team from Sandia National Laboratories (SNL) was given several access points within the architecture to simulate worst-case scenario attack events and to test the intrinsic situational awareness capabilities of the sensor architecture. The SNL cyber red team accessed the PNNL test architecture remotely via a virtual private network (VPN). The sensors were "placed" in several sections of the test architecture to determine if certain types of sensing were more appropriate than others for a utility's IT and OT networks. Using Design of Experiments methodology, we developed an experiment to test the cyber defense-in-depth. Our Fundamental Problem was to determine if a layered, cyber security sensing strategy provided any situational awareness improvement over current standard systems? From this problem came the overall demonstration question:

How useful/successful are intrusion detection sensing strategies, using heterogeneous inspection techniques at various locations on the PERA model, at identifying malicious activity of a knowledgeable and determined cyber adversary?

We developed multiple hypothesis to explore to answer our fundamental problem:

1. Hypothesis #1:

H₀: Aggregate and individual CIA scores on the defense-in-depth utility control system network are equivalent to the corresponding scores observed on the standard utility control system network.

H_{Alternate}: Aggregate and individual CIA scores are significantly different between the standard and the defense-in-depth utility control system networks.

2. Hypothesis #2:

H₀: Intrusion detection at various locations does not improve the situational awareness (i.e., provides non-useable or unmanageable information) for defense of a utility control system network.

H_{Alternate}: Intrusion detection at various locations provides a significantly improved situational awareness for defense of a utility control system network.

3. Hypothesis #2A: Sensor Location Hypothesis

H₀: The sensor mean precision scores on the corporate (Enterprise) network are equivalent to the sensor mean precision scores in the operations network.

H_{Alternate}: The sensor mean precision scores are significantly different between the two networks.

Sensor location was chosen as a factor to determine if cyber intrusion detection sensing is more useful in certain areas of a utility company than others. Utility companies do not necessarily have the expertise or manpower to actively and persistently defend all of their networks, so cyber intrusion detection sensors need to be properly placed for optimal cyber situational awareness.

We defined sensor location as the area of the network from which network traffic (i.e., packet capture) was passively duplicated and sent to the NIDS for analysis. The traffic/packet capture was done at the switch level. All traffic to and from all switch-connected devices was evaluated by the NIDS.

4. Hypothesis #3A: Sensor Type Hypothesis: Corporate Network

H₀: The mean precision scores for each NIDS type on the corporate network are equivalent.

H_{Alternate}: There are at least two NIDS types that have a mean precision score that is statistically significant from each other on the corporate network.

5. Hypothesis #3B: Sensor Type Hypothesis: Operations Network

H₀: The mean precision scores for each NIDS type on the operations network are equivalent.

H_{Alternate}: There are at least two NIDS types that have a mean precision score that is statistically significant from each other on the operations network.

Variables

Independent variables are the manipulated factors in an experiment, where the changes should affect the dependent variables' measurements. The results of an experiment with independent and dependent variables can reveal the relationship (i.e., correlation) between the variables. Controlled variables are held constant to allow the independent variables to be better understood.

Independent Variables:

1. Utility control system network defensive posture
 - A Level A: Standard defense for utility control system
 - B Level B: Defense-in-depth for utility grid
2. Red team access locations
 - A Corporate network
 - B Operations network
3. NIDS Location
 - A Level A: Corporate network
 - B Level B: Operations network
 - C Level C: All above locations

Sensing Type. Sensing type was another independent variable chosen as a factor to determine if certain modalities of cyber intrusion detection sensing are more useful than others in a utility company's networks.

A utility company has proprietary and Personal Identifiable Information to protect, like other companies, but also has the added responsibility of protecting extremely expensive and heavily relied upon field devices (such as boilers and steam turbines). Sensing types each have their strengths and weaknesses, and matching one, some, or all to a utility's specific networks was the goal for this factor. We know from prior experimentation on military C2 networks that heterogeneous sensor combinations can increase attack detection rates while also offering an approach through correlation to controlling false-positive rates.

4. Signature-Based

Signature-based sensing is the cyber defense industry standard for intrusion detection. Signature-based sensors use periodically updated local databases of the specific patterns, such as byte sequences in network traffic of known malicious attacks. Although signature-based sensing can readily detect known attacks, the possibility of detecting new or custom attacks, or polymorphic versions of known attacks is significantly lower.

5. Anomaly-Based

Anomaly-based intrusion detection systems were introduced to cyber defense as a mechanism for catching unknown attacks. The anomaly-based approach creates a baseline model of the network, and then detects any traffic, host, or condition that falls outside the model parameters. Traditionally, this method of intrusion detection suffers from a high false-positive rate as previously unknown legitimate activity, such as a seldom used device, is classified as malicious.

6. Machine Learning-Based

The machine learning-based intrusion detection approach is similar to the anomaly-based approach in that a baseline model of the network is created. The main difference, however, is that a model of attack traffic is also created. Machine learning has at least two labels/categories (known normal and known attack) to compare new traffic against. Machine learning-based intrusion detection has been proven to catch unknown and custom attacks, but also traditionally suffers from a high false-positive rate. There is also very promising evidence that machine-learners can detect polymorphic versions of known attacks.

Dependent Variables

1. Data exchange scores
 - a. PLC-to-SCADA Server
 - b. SCADA Server-to-HMI
 - c. Field Device-to-PI Node
 - d. PI Node-to-PI Server

2. NIDS performance

Sensor Recall and Precision

Intrusion Detection sensing can only provide cyber situational awareness of threats. The quality of that situational awareness was quantitatively assessed by grading the sensor outputs. By analyzing the sensor events/alerts after each test run, we categorized each alert as a true or false-positive based on ground truth provided from the cyber red team on the number and type of attack (as well as their Internet Protocol (IP) addresses and Media Access Control (MAC) addresses). Similarly, for each red team attack, we identified which attacks were caught or missed. These sensor outputs became our measures of performance and the basis for our measures of evaluation.

- a. Recall (From red team attacks (exploits) via red team logs)
 - 1) Caught Attacks
 - 2) Missed Attacks
- b. Precision (From sensor alert Logs (events) via syslog messages to the SIEM)
 - 1) True-Positive Alerts (correctly identified events)
 - 2) False-Positive Alerts (incorrectly identified events)

Using the above data, quantitative assessment of sensor performance was determined by the percentage of correct alerts (i.e., precision) and the percentage of attacks actually caught (i.e., recall).

$$\text{Equation 1} \quad \text{Recall} = \frac{\text{Caught Attacks}}{\text{Missed} + \text{Caught Attacks}} = \frac{\text{Caught Attacks}}{\text{Total Attacks}}$$

$$\text{Equation 2} \quad \text{Precision} = 1 - \frac{\text{FalsePositives}}{\text{TruePositives} + \text{FalsePositives}} = 1 - \frac{\text{FalsePositives}}{\text{TotalPositives}}$$

Sensor recall and precision were calculated for the individual sensor, the sensing inspection method (machine-learning, anomaly or signature-based), and in combination after aggregating at the SIEM-level.

Controlled Variables

The following precautions were taken during testing to control variables (i.e., holding constant the parameters that could inadvertently affect experiment results of dependent variables):

- a. Traffic Generation.

Traffic generation for the environment was provided by PNNL's emulation software suite, Raven, and custom scripts actuating field device components. The Raven traffic and field device scripts were restarted at the beginning of each day and each test run.

- b. Utility Control System Network Configuration and Function.

Once the testing started, no more changes were allowed to the environment to ensure a consistent landscape across the test runs.

c. Blue Team Response Actions.

In exercises where the objective is personnel training, the blue team members (network defenders, utility operations crews, system administrators, etc.) are graded on their responses to situational data. Although the sensors and other individual technologies provided a great deal of insight into the actions of the red team, we did not allow any action or reaction by individuals or technology providers. The objectives of the test were to assess the aligned sensors and the cyber Defense-In-Depth technologies; not the human operators using those devices or running the utility control system.

d. Control and monitor potential scores for data exchanges.

e. Defense-in-depth individual endpoint technologies (i.e., all or none; no individual assessment).

f. Red team knowledge or learning curve.

Test Environment

The test environment was built and run by PNNL and incorporated design input from HECO, Sandia National Lab, and Resurgo SMEs. Within one environment, PNNL operated two utility control systems. These control systems shared certain common equipment such as traffic generators, some ICS equipment, and field device simulators. Ideally, the whole environment would have been simulated and tested, but funding and time constraints narrowed the focus of testing to the most pertinent issues. A conceptual diagram of a utility's network infrastructure can be seen in the Figure below. The grey areas in the diagram are network areas de-emphasized for the test due to a lack of time and resources. Also, giving the red team direct access to the operations networks reduced the need for the de-emphasized network sections. Normally, for hackers to attack utility grid operations from the Internet, they would first need to establish a beachhead in the operations or corporate networks. We gave the cyber red team access to these networks to save experiment time and resources, so establishing a beachhead (via phishing for example) was not necessary and thus negated the need to fully simulate all of a company's networks and access points.

5.1.1 Level A: Enterprise Network

The Enterprise or corporate network contained enterprise-level systems (e.g., mail server, domain controller, etc.) that are common within utility company networks. Packet capture of the corporate network included the Raven generated traffic, PI server traffic, and any Raven "agent" traffic. The Raven traffic generation provided the sensors with realistic traffic data to analyze to include user agents accessing a simulated Internet. This type of traffic and sensing location represented a normal large business work environment.




Figure 10. Depiction of Test Environment

5.1.2 Level B: Operations Network

The operations / operations technology network(s) contained all the SCADA equipment required to monitor and control the utility field devices. The basic devices within a typical operations network are the HMI, SCADA server, PLC, and data historian (although equipment providers name their company/proprietary products and systems differently). The functional type of traffic associated with this network varied greatly in size, format, and complexity from the relatively diverse corporate network.

The operations networks traffic also contained Raven generated traffic, although there is some question as to how appropriate this was for the test architecture OT networks. The randomness of traffic generation is not representative of actual traffic on OT networks and presented challenges for the machine learning and anomaly detection sensors.

Network Architecture

The network architecture implemented at PNNL is shown below, where the “operations” network was split into the two treatments: Defense-in-Depth (on the left) and NIST-compliant (on the right).




Figure 11. Depiction of The PNNL Test Environment

Constraints, Limitations, and Inconsistencies

As with any test or experiment, the project had to be conducted within given constraints or limitations to ensure our objectives could be met. Although constraints on time, funding, access to venues, available equipment etc. are typically viewed negatively, they also have positive influences. Operating within constraints encourages novel solutions and approaches. In other words, constraints and limitations are not necessarily negative because they can encourage ingenuity and creative thinking. The following were the major constraints that we had to address in our planning, limiting our testing options:

- **Red Team Funding**

The red team was not funded adequately to allow for onsite participation at PNNL. Their remote participation was sufficient to produce experiment results, but onsite participation would have resulted in a more comprehensive validation of attacks, attack effects, and test data.

- **Experiment Length**

Due to travel and personnel costs, the onsite time at the PNNL test venue was limited to two weeks. This restricted the number of test runs and test samples for the experiment. Effects are harder to detect in smaller sample sizes. A greater number of test samples would have allowed for more statistical power in the analysis.

- **Experiment Factors**

With more time and money, experimentation on more independent and dependent variables would have been possible. Many of the controlled (i.e., held constant) variables could have been tested, while also exploring of human and subjective results.

- **Projection of Sensor Results**

The quantitative performance results of the sensors cannot be directly attributed to performance on a real utility's networks. The use of traffic generation, which can be noisy and inconsistent, caused the sensors (especially anomaly-based) to experience more total alerts and false-positives than normal. On the other hand, the lack of experiment time required the red team to attack the networks in a "loud and obvious" manner. This probably resulted in higher percentage of attacks caught by the sensors. In the next event in the series, fault and intrusion tolerant defense with associated sensing strategy was employed on a functional HECO grid without artificial traffic generation. This provided a relatively quiet yet realistic environment to assess sensor performance.

Cyber Red Team Attack Variance

Because the red team was not following an attack script, the quantity and duration of attacks varied greatly from run to run.

The red team was permitted to fight freely with no time constraints or script requirements because the various segments of the architecture contained different devices and the red team had no chance to work with the sensors or the cyber defense-in-depth technologies in advance.

5.1.3 Baseline Characterization

Multiple network traffic baselines were required to prepare the cyber defense-in-depth technologies for installation at PNNL.

Network traffic for the control system was captured prior to our arrival and mailed to Resurgo to train the machine-learning NIDS. Resurgo copied the traffic and provided it to Imprimis to permit them to design and write the rules for the CTAM-R NIDS. Resurgo also provided a copy to JHU so they could ensure their systems were compatible with the control system network.

On arrival at PNNL, network traffic was captured again to determine if the traffic had changed. Network traffic was captured a third time once all of the cyber defense-in-depth technologies were installed, tuned and operational. This was done to permit the NIDS to be trained/familiarized with the effects the other technologies had on the control system network. By the end of the first week, we knew what the original network baseline traffic was and had a second baseline with the technologies installed.

Simultaneously, the white cell used the first week at PNNL to familiarize themselves with the normal, baseline function of the data exchanges and the corresponding physical and emulated ICS devices (e.g., turbine and boilers, their control commands and their output responses).

5.1.4 Design and Layout of Technology Components

A diagram of the machine-in-the-loop emulation of the utility control system network can be seen below. The “Fake Internet” represented External Connections to the commercial power provider or utility, the Enterprise represented the administrative functions of the power provider, Information Assurance was an out of band network fed by one-way network taps from the Enterprise and Operations networks, and the Operations Network was where the undefended network segment and the defensive technologies defended segment were installed the test environment. The Sandia National Labs cyber red team was given access directly to the Enterprise Network at Red Team Access Point 1. Access here simulates the red team having successfully penetrated the Enterprise Network by social engineering, phishing or other means. This was done to save time and focus the red team on the data exchanges and technologies under test vs spending the limited time available attempting to gain access to the network. Time constraints to the project narrowed the focus of testing to the most pertinent issues.

Data Collection Equipment: A SIEM in the out of band network collected all alerts and events from the technologies and NIDS. PNNL instrumented the entire baseline network to track network communication using an open-source application, EtherApe. EtherApe is a graphical network traffic display tool that views traffic end to end and supports a variety of frame and packet types. During testing EtherApe was only used as a white cell situational awareness tool.




Figure 12. System Depiction

5.1.5 Operational Testing

The first week of the PNNL experiment focused on the setup and integration of the defense technologies into the network architecture. Care was taken to ensure that technology integration did not affect the baseline function of the architecture:

- A. The Information Assurance network and associated IDS equipment was placed in an out-of-band network. This network was in a passive mode and only received data. No other interaction was possible.
- B. The defense-in-depth technologies in the Operations network were contained in a separate network segment (via Virtual Local Area Network (VLAN)). The technologies were not able to interact with the baseline architecture in a separate VLAN. The figure above, depicts the segmented networks as VLANs 200 (defended) and 250 (baseline).

While the primary goal for the first week was to integrate all the cyber defense-in-depth technologies into the test environment, the red team used that time to familiarize itself with the test architecture and to observe the installation and actions of the defense technologies. In coordination with the red team, network traffic was captured when they were not in the network. The traffic was then cleaned, prepared, and used to train the machine-learning NIDS. Once all technologies were installed and tuned, a dry run was conducted using previously recorded threat traffic. This ensured that the NIDS functioned and offered the possibility of causing the Prime and Spines technologies to react.

The second week at PNNL focused on testing. Using a separate VLAN for the baseline and defended network segments required placing the test environment into specific configurations for each test run. This was done by connecting or disconnecting VLAN 200 or VLAN 250 from the architecture as required so that only the network segment under test was connected at any one time. Test runs started with PNNL and the technology providers reporting ready. The white cell then called Sandia who used a Virtual Private Network (VPN) to connect to Red Team Access Point 1 or 2 as directed by the white team. There were four testing days and four test runs the second week. Each test run was given one day. Test runs were not limited to a set amount of time. The red team was given until the end of their business day to attack the designated communications exchanges and defending technologies. Some security postures provided little attack surface for the red team, so test time was better utilized by the red team reporting that it could find no access and requesting to move to the next run or for conditions to be changed to permit the run to continue.

Table 6. Timeline of PNNL Test Activities (March-April 2017)

	27-Mar	28-Mar	29-Mar	30-Mar	31-Mar	01-Apr	02-Apr	03-Apr	04-Apr	05-Apr	06-Apr	07-Apr
1st Week												
Inventory Equipment. Begin set up												
Connect Technologies to PNNL network												
Connect Technologies to Test network												
Run data over PNNL network.												
- validate technology performance												
- validate technology compatibility												
Run attacks over PNNL Network												
2nd week												
Validate technology / tune / repair												
1 x AM & 1 x PM iterations with Red team												
Evaluate data and defense performance												
1 x AM & 1 x PM iterations with Red team												
Equipment breakdown and shipping prep												

All equipment was removed and taken home by the technology providers.

5.1.6 Sampling Protocol

No samples were collected. All network traffic during all test runs was captured and stored for analysis. The span ports sent 100% of the network traffic from all segments to data storage. This permitted forensic analysis on return to home station to permit initial results to be challenged and validated.

- *Calibration of equipment.* As the technologies were installed they were configured and optimized for the PNNL network. Data capture equipment required no calibration.
- *Quality assurance sampling.* Other than the dry test run the first week to verify defense technologies and data capture equipment performance, no quality assurance sampling was conducted.

5.1.7 Sampling Results

All network traffic during all test runs was captured and stored for analysis. The span ports sent 100% of the network traffic to data storage. Each test run was monitored and the final data file checked to ensure data capture.

5.2 HECO OPERATIONAL VALIDATION CONCEPTUAL DESIGN

Activities at HECO compared the company's normal control system network architecture (baseline) to a second, similar, portion of the control system architecture with the cyber defense-in-depth technologies installed. HECO's architecture exceeds the Purdue Enterprise Reference Architecture (PERA) model and NIST recommended cyber security implementations. The plant also complies with the NERC-CIP's nine standards and 45 requirements covering the security of electronic perimeters and the protection of critical cyber assets as well as personnel and training, security management and disaster recovery planning.

Some of the NIST and NERC-CIP standards and requirements were "waived" in order to get the Certified Ethical Hacker (CEH) red team inside the power plant to permit them to focus on testing the cyber defense-in-depth technologies. The red team was given access points within both architectures to simulate an already compromised network (worst-case scenario) and to test the intrinsic situational awareness, intrusion tolerance, and "fight through" capabilities of the two architectures.

Fundamental Problem: Does a layered, aware cyber security defense-in-depth affect the performance of a power plant?

Demonstration Question #1: How does a layered, aware cyber security defense-in-depth affect the latency and traffic of a control system in a functional power plant compared to the current state-of-the-art cyber defense for the control system within the plant?

Demonstration Question #2: Can a machine-learning intrusion detection system be successful in detecting a knowledgeable and determined adversary in the control system network of a functional power plant?

5.2.1 Conceptual Test Design

We developed a simple pass / fail test to determine which hypothesis was correct.

A. Hypothesis #1:

H₀: The installed cyber defense-in-depth technologies induce latency, degrade performance or cause devices to perform outside of specified or accepted tolerances compared to the corresponding performance of the baseline standard utility control system network.

H_{Alternate}: The installed cyber defense-in-depth technologies do not induce latency or affect device or system performance of the utility control system network.

B. Hypothesis #2:

H₀: Machine-learning intrusion detection does not improve the cyber situational awareness (i.e., provides non-useable or unmanageable information) for cyber defense of a utility control system network equivalent to the corresponding detection rates for the standard utility control system network's signature and anomaly NIDS.

H_{Alternate}: Machine-learning intrusion detection provides significantly improved situational awareness for cyber defense of a utility control system network compared to a standard utility control system's defenses.

C. Independent Variables:

1. Utility control system network defensive posture
 - Level A: Standard defense for utility control system
 - Level B: Defense-in-depth for utility grid
2. Red team access locations were only within the control system / OT network. The HECO corporate / Enterprise network was off limits and not involved in testing.

D. Dependent Variables:

1. Data latency (as measured by HECO)
2. NIDS performance

E. Controlled Variables:

1. Traffic generation
2. Control system network configuration and function
3. Defense-in-depth individual endpoint technologies (i.e., all or none; no individual assessment)
4. Red team knowledge or learning curve

F. Test Design:

The test environment was a power plant provided by HECO modified to permit portions of a control system to be independently tested. The plant had all normal HECO instrumentation as well as additional instrumentation added by Resurgo and Johns

Hopkins. Within the power plant, one section of the control system was augmented with the defense-in-depth technologies and a second section remained as HECO normally operated that control system.

The Enterprise Network, all external connections, and the Enterprise DMZ were not part of the demonstration. The External Connections and the Enterprise DMZ were de-emphasized since the red team was physically within the power plant. This saved demonstration time and focused the red team on the devices and technologies under test vs time they would have spent attempting to gain access to the plant's network. The Enterprise Network was off limits to the red team to permit normal HECO operations to continue unimpeded.

5.2.2 Baseline Characterization

HECO collected control system traffic three weeks prior to the demonstration start to establish a system baseline for MANA™ and Spire network traffic familiarization. Traffic was collected again after all Spire technologies were installed. This provided a true baseline and a defended baseline for comparison and analysis of results.

5.2.3 Design and Layout of Technology Components

The testing at PNNL established that NIDS location did not affect any of the other variables (due to the use of passive network taps). For the HECO event, the red team was physically within the power plant to meet HECO security requirements. The configuration depicted in the figure below shows how the network and equipment were set up for the HECO demonstration.

Within the power plant one section of the control system was augmented with the defense-in-depth technologies and a second section remained as HECO normally operated that control system.

NOTE: The Enterprise Network, Enterprise Firewall, the Enterprise DMZ, the Operations Firewall, and the Operations DMZ were not part of the demonstration. Because the red team was physically within the power plant and HECO was conducting normal business operations, all of these areas were eliminated from participation in the demonstration.




Figure 13. System Depiction with Defense-in-Depth Components

This saved demonstration time and focused the red team on the devices and technologies under test in the plant vs time the red team would have spent attempting to gain access to the plant's network.

- Major Components of the System (from top to bottom in above diagram):

1. The Enterprise Firewall. This firewall separates the Enterprise network from the Operations Network. This is a business best practice and a PERA recommendation to separate the business (IT) network from the OT network. The intent of the firewall is

- to prevent attacks originating in the internet from penetrating thru to the OT network. This was not part of the demonstration.
2. The Enterprise Network. This is where HECO conducts commercial business and was not part of the demonstration.
 3. The Enterprise DMZ. Used in conjunction with a firewall, its purpose is to highlight traffic crossing a “no man’s land” separating the Enterprise network from the OT network. This was not part of the demonstration.
 4. The Operations Firewall (Access to Enterprise and Operations Networks): Standard utility cyber defense places a firewall between the corporate and operations networks. Here HECO exceeds the PERA standard and adds a second firewall to make it even more difficult for a threat coming in from the internet thru the enterprise network to reach the OT network. This was not part of the demonstration.
 5. The Operations DMZ. Used in conjunction with a firewall its purpose is to highlight traffic crossing a “no man’s land” separating the Enterprise network from the OT network. This was not part of the demonstration.
 6. The DCS Switch 3 created the Management LAN which contained the IDS and Spire equipment. This LAN was an out-of-band network that was in a passive mode and could only receive data. No other interaction was possible with this LAN segment or the devices on it except by directly, physically, plugging in to a device.
 7. The DCS Switch 1 was part of the Operations Network contained the Spire components as well as the HECO workstations that controlled and monitored the field devices. Contained within this network were:
 - a. PI [Interface] Node: the PI node contains the interfaces that pull data from the various field devices or controllers. The data is then sent to the PI server in the corporate network.
 - b. Human Machine Interface (HMI): the control and monitoring station used to run the daily operation of the field devices. HMI hardware can vary greatly depending on function, and the client software is usually proprietary based on the type of controller (e.g., Siemens, ABB, etc.). The HMI software is designed to act as the main user interface while also providing enhanced graphing and display capabilities.
 - c. [Process] Control Server: While the HMI may be the actual user terminal, the control server functions as the workhorse for collecting/saving data and issuing the actual control commands to the controllers.
 - d. Energy Management System (EMS): these systems manage energy production distribution and keep power production smooth over a utility’s multiple power plants.
 - e. PLC: controllers are ruggedized computers made specifically to function within the demanding tolerances of operational environments. PLCs can be programmed to use a variety of inputs and outputs, and in many cases utilize Ethernet at the edge of the operations network.

8. Field Device Loop: this network housed the actual equipment of interest and consequence, that is, the equipment that would be targeted by advanced persistent threats. Destruction or damage to the field devices would cause significant costs on many levels. All HECO field devices will be disconnected from the two control systems under test to prevent damage or spurious outputs.
 - a. Turbine: utility control system equipment (simulated at PNNL) used for producing continuous power by rotating a blade assembly with water, steam, gas, air, or other medium.
 - b. Boiler: utility control system equipment (simulated at PNNL) that heats water to generate pressurized steam, which in turn activates a turbine.

5.2.4 Operational Testing

The white cell used the first week at HECO to familiarize themselves with the normal function of the plant's networks, control systems and equipment. The plant was offline (not producing power) and had limited variation in its control system communication, so the familiarization process did not take long. In preparation for data collection, a SIEM was installed to collect all alerts and events from the control system equipment, PLCs and Spire technologies. HECO had its normal instrumentation in place across the entire network to track network communications and observe red team actions.

The primary goal of the first week of the HECO demonstration was to integrate all the cyber defense-in-depth technologies into the plant in readiness for testing. Simultaneously, the CEH "red team" used that time to familiarize themselves with the test architecture and observe the installation and actions of the defense technologies. This made the red team as knowledgeable as possible about the plant and its systems to simulate an APT or other nation-state threat that has had time to research, conduct reconnaissance and establish a foothold. These are the normal first steps in conducting an attack and constitute the beginning of the threat's "cyber kill chain". All of the technologies (MANATM NIDS, SIEM and Spire control and monitoring equipment) were placed behind DCS Switch 3, in an out-of-band network. This network was in a passive mode and only received data. In coordination with the red team, network traffic was captured to train the MANATM when the red team was not in the network. The traffic was then cleaned, prepared, and used to train the NIDS to the plant network with the Spire technologies in place and communicating. Once all technologies were installed and tuned, a dry run was conducted using with the red team provided attacks. This both verified that the technologies were functioning and allowed the red team to test some of their tools and assess their success.

The second week at HECO focused on operational validation. Four days of eight-hour operation were planned, but due to interruptions by over 100 visitors ranging from Senate senior staffers to engineers from mainland cooperatives, our testing became a series of two-hour iterations, repeated multiple times a day. These further evolved into a demonstration on two large screens observable by the visitors. On one screen, the red team used the exploits that worked best in an attack demonstration that started with reconnaissance and concluded with them overriding system safety settings and blocking operator access to the safety overrides so that they could not be reset. On the other screen, the visitors could observe the MANA's correlated high confidence alerts as they were generated and displayed in real time. From there, visitors were shown the

Spire control system and then put into safety gear and taken into the plant to see what had been attacked, where the plant operators would monitor the defense in depth technologies when installed, and where Spire's components were installed.

A. CEH Preparation

To make sure the red team was adequately prepared for the event; all test environment information was shared with them beforehand. They were also given control system network traffic and source code for the defense technologies (minus the NIDS). An additional benefit to this approach was to reduce the learning curve for the red team, as we did not want them to discover important new information halfway through testing. Learning significant new information during testing makes it hard to compare test runs since it affects how the red team performs. The red team was the impetus to the determining NIDS and defensive technologies performance in the power plant (dependent variable). We needed a consistent effort from the red team for every test run.

B. System Checks (Week One)

Upon arriving at HECO, the first week was dedicated to equipment setup and permitting the red team to continue to practice and conduct reconnaissance (familiarize themselves) with the network. Again, this effort reduced the learning curve for the red team; however, the primary goal for this week was to integrate all the defense-in-depth technologies into the power plant and to permit HECO engineers to verify that the technologies did NOT affected plant control system or device performance.

C. Test Runs (Week Two)

The first day of the second week was used to perform the first test runs. However, due to interruptions by the first visitors, Day 1 began our evolution into our testing becoming a series of two-hour iterations, repeated multiple times a day. We learned from each interruption and iteration and developed a repeatable demonstration observable by the visitors.

An additional benefit of red team preparation/familiarization and the interruptions was that we did not need a second replicate of test runs. The repetitive demonstrations essentially provided us with test replication.

Table 7. Timeline of the HECO Demonstration (Jan-Feb 2018)

	22-Jan	23-Jan	24-Jan	25-Jan	26-Jan	27-Jan	28-Jan	29-Jan	30-Jan	31-Jan	01-Feb	02-Feb
1st Week												
Inventory Equipment. Begin set up												
Connect Technologies to HECO network												
Connect Technologies to Test network												
Run data over HECO network.												
- validate technology performance												
- validate technology compatibility												
Run attacks over HECO Network												
2nd week												
Validate technology / tune / repair												
One test run per day												
Evaluate data and defense performance												
VIP Briefing (Technology transition activity)												
Excursion (If NAVFAC / HECO / Red request)												
Equipment breakdown and shipping prep												

5.2.5 Sampling Protocol

Data Collection (Week Two):

Data was collected automatically and manually during each test run. During and after each test run, the white cell assessed the results manually.

Automatic data collection occurred during all test runs. All logs and alerts were sent to the SIEM, a syslog server, as is normally done for Cyber Network Defense.

Data Collector(s): All results were displayed in near real time on large monitors. HECO and the SMEs assessed attack/defense effectiveness.

The Resurgo engineers operating the syslog server were the data collectors. All runs and results were saved on the syslog server.

Data Recording: The actual data recording was done automatically via the syslog server every test run. The syslog server also acted as the SIEM. All data sent to the SIEM was manually downloaded to an additional storage media for later analysis during the demonstration. NOTE: at HECO's request all data was all wiped to protect information about plant systems at the conclusion of the experiment.

HECO's normal instrumentation was operational and collected data on both network segments for HECO internal consumption. HECO provided observations and notes from its instrumentation. See note at chapter end.

Data Description: Data was a mixture of machine logs, visually displayed alerts, and SME assessments.

Data Storage and Backup: All packet capture of the network traffic was saved by the NIDS. However, at HECO's request this data was wiped to protect plant system information. See note at chapter end.

5.2.6 Sampling Results

No sampling was done during the operational validation.

NOTE: Removal of any data by Resurgo, or the other technology providers, would have been a violation of the NERC rules in accordance with the Federal Energy Regulatory Commission's rules, regulations, and orders, as well as NERC's Rules of Procedure, Appendix 4C (NERC Compliance Monitoring and Enforcement Program (CMEP)) and violations of the Critical Infrastructure Protection (CIP) NERC Reliability Standards for Critical Cyber Assets (CCAs). Data removal poses a serious or substantial risk to the reliability of the bulk power system (BPS). Exposure of the username and cryptographic information concerning a power plant could aid a malicious attacker in using this information to decode the passwords and increases the risk of a malicious attacker gaining both physical and remote access to systems.

Page Intentionally Left Blank

6.0 PERFORMANCE ASSESSMENT

6.1 Phase 1 - Risk Reduction Demonstration (PNNL) Data Analysis

6.1.1 Performance Objective 1 – Architecture Cyber Security

Table 8. Performance Objective One

Performance Objective	Metric	Data Requirements	Success Criteria
Architecture Cyber Security	Data Stream Protection Score	Confidentiality Score	Contents unknown to Red Team
		Integrity Score	Contents unchanged by Red Team
		Availability Score	Contents unhindered by Red Team

The “white cell” gave a score to each data exchange relative to how the data exchange withstood red team attacks during test runs. Each individual data exchange score was aggregated to create a value representative of the whole system (i.e., cyber-secure utility architecture). There are two responsibilities for the SCADA data exchanges: monitor and control. The minimum score for the architecture is 0 (if the red team violates all confidentiality, integrity and availability requirements) and the maximum score is 13 (if the red team affects nothing).

Table 9. CIA Scoring Example

Information Assurance Requirements	Monitor	Control
Confidentiality	<i>Low(1)</i>	Medium(2)
Integrity	<i>Medium(2)</i>	High(3)
Availability	<i>Medium(2)</i>	High(3)
Maximum Score	5	8

To determine the statistical significance of deploying a Defense-in-Depth architecture versus the standard NIST-compliant architecture and to determine if the red team access location affected results, we used the Paired T-Test. Since the standard NIST-compliant and Defense-in-Depth architectures differ only in the addition of the defense-in-depth technologies, we did not consider the two to be independent from each other. As such, we used a statistical test that considers one subject under two different treatments (i.e., paired T-test). The scores compared for significances are defined by the Sandia National Laboratory Reference Architecture recommended CIA (Confidentiality, Integrity, and Availability) criteria.

The red team was tasked to determine if they could disrupt communications by exposing, denying, or modifying specified data exchanges. The value of a truly secured architecture is defined by the protection of its data exchanges.

6.1.2 Performance Objective 2 – NIDS Performance

6.1.2.1 Alert Recall Analysis

Table 10. NIDS Recall Performance Objective

Performance Objective	Metric	Data Requirements	Success Criteria
Sensor Performance	Sensor Recall	Caught Attack	Correctly identified Red Team attack
		Missed Attack	Did not fail to identify a Red Team attack

Alert recall was calculated with “ground truth” knowledge of the number and type of attacks provided by the SNL cyber red team in a master time-line log after each test run. The log contained the malicious activity with a recorded date/time of the activity. A total of 53 attacks were in the master time-line log and they are listed in Appendix A. The master time-line also contained the target IP addresses that SNL exploited. Resurgo engineers used the target IP addresses, known SNL IP addresses, and dates/times to determine if the malicious activity was caught and reported by MANA™ or CTAM-R by reviewing the NIDS output logs.

To account for any discrepancies in the reported time in the master time-line and sensors’ output logs, Resurgo LLC engineers applied a time delta when parsing through the logs to allow for a reasonable time window. The time window, T , provided the necessary window to validate red team caught-attacks with sensors output alert logs and is detailed in Equation 3.

$$\text{Equation 3} \quad T: t_a - \Delta t \leq t_{alert} \leq t_a + \Delta t$$

The time delta was set at $\Delta t = 600 \text{ sec}$ or 10 minutes, where t_a is the malicious activity time recorded by SNL, Δt is the delta time, and t_{alert} is the recorded time from a sensor alert file. For an attack to be graded as “caught”, a sensor alert log had to contain a red team IP address and the target IP address within the defined time window T . The recall was then calculated as defined in Equation 1.

$$\text{Equation 1} \quad \text{Recall} = \frac{\text{Caught Attacks}}{\text{Missed} + \text{Caught Attacks}} = \frac{\text{Caught Attacks}}{\text{Total Attacks}}$$

6.1.2.2 Alert Precision Analysis

Table 11. NIDS Precision Performance Objective

Performance Objective	Metric	Data Requirements	Success Criteria
Sensor Performance	Sensor Precision	True-Positive Alert	Correctly identified Red Team attack
		False-Positive Alert	Did not categorize normal traffic as Red Team activity

The process for determining alert/log ground truths was much different than the recall/attack analysis. Each individual alert log had to be validated as a true or false-positive before alert precision could be calculated. Most alerts were validated from the set of known and determined red team IP addresses; alert logs containing a red team IP address were automatically graded as valid alerts (i.e., true-positives).

However, simple IP address assessment was not enough to validate the 300,000+ alerts, since some attack behavior manifested in the alerts with non-red team IP addresses. For instance, alerts caused by Man-in-the-Middle (MitM) attacks contained normal host-network source and destination IP addresses.

To account for these more complex attacks, Resurgo, LLC, engineers did additional packet analysis to identify combinations of red team MAC addresses with normal host-network IP addresses. Any sensor log containing a known red team MAC address was labeled as a valid attack alert (i.e., true-positive).

For sensor logs not containing MAC addresses, any sensor log generated during the attack timeframe (identified by red team MAC address information) was also graded as valid true-positive alert. The Figure below details the process of labeling each alert from a sensor output.




Figure 14. Alert Analysis Flow Chart

After labeling each log file, calculating the alert precision of each sensor was possible using Equation 2.

$$\text{Equation 2} \quad \text{Precision} = 1 - \frac{\text{FalsePositives}}{\text{TruePositives} + \text{FalsePositives}} = 1 - \frac{\text{FalsePositives}}{\text{TotalPositives}}$$

NIDS recall and precision was calculated for each individual NIDS and after combining both NIDS results at the SIEM-level. The difference in results demonstrated the effects of implementing the heterogeneous NIDS.

$$\text{Recall} = \frac{\text{Caught Attacks}}{\text{Missed} + \text{Caught Attacks}} = \frac{\text{Caught Attacks}}{\text{Total Attacks}}$$

$$\text{Precision} = 1 - \frac{\text{FalsePositives}}{\text{TruePositives} + \text{FalsePositives}} = 1 - \frac{\text{FalsePositives}}{\text{TotalPositives}}$$

No analytical methodology was used for comparison with the control group (standard utility architecture) since intrusion detection is not a standard practice within control system networks.

6.1.3 Performance Assessment

We looked at each individual criterion separately to determine if deploying the Defense-in-Depth architecture versus the standard NIST-compliant architecture was statistically significant for each criterion. More importantly, we also looked at the system as a whole to analyze the aggregate score of each data exchange.

Shapiro-Wilk (Normality) Test for NIST vs Defense-in-Depth

Prior to running the tests, we checked that the assumptions of the statistical tests were met, as the sample data must come from a normal distribution. We tested for normality using the Shapiro-Wilk test. The Shapiro-Wilk test assesses the null hypothesis that the scores were drawn from a normal distribution. The results of each CIA criterion and the Aggregate scores are detailed below.

H₀: The scores were drawn from a normal distribution.

H_{Alternate}: The scores were not drawn from a normal distribution.

The significance level for this test was chosen to be $\alpha = 0.05$, where p-values above this failed to reject the null hypothesis (i.e., the data is indeed from a normal distribution).

Table 12. Architecture Type Aggregate Shapiro-Wilk Test

Aggregate Shapiro-Wilk Test	DF	Test Statistic	p-value
Sample Data	5	0. 820169	0. 088550

Table 13. Architecture Type Confidentiality Shapiro-Wilk Test

Confidentiality Shapiro-Wilk Test	DF	Test Statistic	p-value
Sample Data	5	0. 821616	0. 091135

Table 14. Architecture Type Integrity Shapiro-Wilk Test

Integrity Shapiro-Wilk Test	DF	Test Statistic	p-value
Sample Data	5	0. 841965	0. 135334

Table 15. Architecture Type Availability Shapiro-Wilk Test

Availability Shapiro-Wilk Test	DF	Test Statistic	p-value
Sample Data	5	0. 775164	0. 034730

Pink cells indicate that the test was not successful, green cells indicate successful passing of criteria.

At the significance level of $\alpha = 0.05$, we determined that our small sample set ($N = 6$) displays a statistically significant difference to a normal distribution in only one (Table 12) of the four scenarios above.

However, non-normality can be difficult to detect in a small sample set. Since T-Tests are quite “robust” to violations of normality, the distribution may be “approximately” normal and still provide valid results. In our case, if we could have increased the number of data exchanges important to the communications of an ICS network, increasing our number of samples, the t-statistic would have converged in probability to the standard (approximately) normal distribution by the law of large numbers. Therefore, based on these results we can assume that our sample data distribution is approximately normal.

Under the assumption that our data is approximately normal, or rather that non-normality is difficult to accurately detect for a small sample set size, we measured the statistical significance of deploying a Defense-in-Depth architecture versus the standard NIST-compliant architecture using the Paired T-Test. We also conducted the non-parametric equivalent to the Paired T-Test, the Wilcoxon Sign-Rank test. The Wilcoxon Sign-Rank test assesses the null hypothesis that two related paired samples come from the same distribution, without the requirement of normality.

6.1.3.1 Paired T-Test for NIST vs Defense-in-Depth

The basic statistics for each criterion is detailed in Table 16 (below).

Table 16. Architecture Type Basic Statistics for CIA and Aggregate Score

Paired T-Test (N = 6) for Architecture Defense Type versus CIA								
	Conf. Mean	Conf. Std. Dev.	Integrity Mean	Integrity Std. Dev.	Avail. Mean	Avail. Std. Dev.	Aggregate Mean	Aggregate Std. Dev
Standard	0.6667	1.4907	1.1667	1.8634	0.3333	0.7454	2.1667	4.0173
Defense-in-Depth	2.000	1.6320	3.8333	1.8634	2.000	1.6320	7.8333	4.5613

Conf = Confidentiality

Avail = Availability

Std Dev = Standard Deviation

We proceeded to test the hypothesis proposed in Section 3.1 Hypotheses:

H₀: Aggregate and individual Confidentiality, Integrity, and Availability (CIA) scores on the Defense-in-Depth utility control system network are equivalent to the corresponding scores observed on the standard utility control system network.

H_{Alternate}: Aggregate and individual CIA scores are significantly different between the standard and the Defense-in-Depth utility control system network.

The significance level for this test is chosen to be $\alpha = 0.05$; where the p-values below this rejected the null hypothesis (i.e., the differentiation in scores between architectures was significant).

The Paired T-Test results for the Aggregate score and each CIA criterion are detailed below in Tables 17, 18, 19, and 20, respectively:

Table 17. Architecture Type Aggregate Paired T-Test

Aggregate Paired T-Test	DF	Test Statistic	p-value
Sample Data	5	-2.6294	0.0466

Table 18. Architecture Type Confidentiality Paired T-Test

Confidentiality Paired T-Test	DF	Test Statistic	p-value
Sample Data	5	-2.0	0.10194

Table 19. Architecture Type Integrity Paired T-Test

Integrity Paired T-Test	DF	Test Statistic	p-value
Sample Data	5	-2.90191	0.03372

Table 20. Architecture Type Availability Paired T-Test

Availability Paired T-Test	DF	Test Statistic	p-value
Sample Data	5	-2.0761	0.0925

Again, pink cells indicate that the test was not successful, green cells indicate successful passing of criteria

After analysis of the Aggregate scores, we note that the overall security of the system shows a significant difference in deploying a Defense-in-Depth strategy. Of the individual CIA components of the data exchanges, only Integrity showed a significant difference between Defense-in-Depth and the NIST-compliant. Confidentiality and Availability failed to exhibit enough difference between scores with the small data size, but are approaching significance.

For brevity to more directly address the overall security of the network, we will continue by analyzing the aggregate scores only from this point forward.

6.1.3.2 Small Sample Wilcoxon Sign-Rank Test for NIST vs Defense-in-Depth

Due to the small sample set ($N = 6$) and the normality assumption being approximate, we also conducted a version of Wilcoxon test where normality is not required.

The W test statistic is calculated as the sum of **all** the positive and negative Sign-Ranks combined (Frederick Mosteller & Robert E. K. Rourke, *Sturdy Statistics: Nonparametrics and Order Statistics*, Addison-Wesley, 1973, 89ff).

H_0 : The median difference in Aggregate and individual Confidentiality, Integrity, and Availability (CIA) scores on the Defense-in-Depth utility control system network and the corresponding scores observed on the standard (NIST) utility control system network is zero.

$H_{Alternate}$: The median difference in Aggregate and individual Confidentiality, Integrity, and Availability (CIA) scores on the Defense-in-Depth utility control system network and the corresponding scores observed on the standard (NIST) utility control system network is not zero.

The significance level for this test is chosen to be $\alpha = 0.05$.

The Wilcoxon Sign-Rank Test results for the Aggregate score are detailed below in Table 21:

Table 21. Architecture Type Aggregate Wilcoxon Sign-Rank Test

Pairs	NIST Aggregate Score: X_a	Defense in Depth Aggregate Score: X_b	S/R of $ X_b - X_a $
1	11	11	---
2	0	11	+3.5
3	0	0	---
4	2	11	+2
5	0	11	+3.5
6	0	3	+1
Test Statistic $\mp W = 10$, $N_{s/r} = 4$			

The $\mp W$ test statistic is the sum of all the Sign-Rank values for the Aggregate score. The $\mp W$ is then compared to critical value found in the table located in Appendix C. The critical value is selected by cross referencing the number of samples and the significance level chosen. In our case, we are performing a directional test at a significance level of 5%. The number of samples that make up the $\mp W$ test statistic is denoted as $N_{s/r}$.

The number of samples for each case after determining the $\mp W$ statistics are all below $N_{s/r} = 5$. No statistically significant conclusion can be drawn with a sample size smaller than $N_{s/r} = 5$. We can however note that the Sign-Rank value of each instance for each criterion was always one-sided (positive). A Wilcoxon Sign-Rank test that has a low number of samples, but displays all the Sign-Rank values one-sided suggests that a larger number of samples would result in statistically significant results. This result comes from the fact that the absolute value of $\mp W$ would only increase given more samples where all the Sign-Rank values are one sided.

By the experiment design, given a larger sample size would have resulted in evidence that the medians of the two samples sets are different. In our case, the Defense-in-Depth scores would have displayed a larger one-sided absolute value of the Wilcoxon Sign-Rank Test Statistic.

6.1.3.3 Shapiro-Wilk (Normality) Test for Red Team Access Location

Prior to running the same tests, with respect to the cyber Red Team's Access Location as the independent variable, we checked that the assumptions of the statistical tests were met. We again tested for normality using the Shapiro-Wilk test. The result of the Aggregate score normality test is detailed below:

H₀: The Aggregate Scores was drawn from a normal distribution.

H_{Alternate}: The Aggregate Scores was not drawn from a normal distribution.

The significance level for this test is chosen to be $\alpha = 0.05$, where p-values above this failed to reject the null hypothesis (i.e., the data is indeed from a normal distribution)

Table 22. Red Team Access Location Aggregate Shapiro-Wilk Test

Aggregate Shapiro-Wilk Test	DF	Test Statistic	p-value
Sample Data	5	0. 708706	0. 007626

At the significance level of 5%, $\alpha = 0.05$, we determined that our small sample set (N = 6) displayed a statistically significant difference to a normal distribution for the aggregate scores. Again, non-normality can be difficult to detect in a small sample set. That is, under the assumption that the distribution will be approximately normal. Since T-Tests are quite "robust" to violations of normality, the distribution may be "approximately" normal and still provide valid results.

Under the assumption that our data is approximately normal, or rather that non-normality is accurately undetectable using the Shapiro-Wilk Test due to a small sample set size, we measured the statistical significance of the Red Team's access locations using the Paired T-Test.

6.1.3.4 Paired T-Test for Red Team Access Locations

The cyber red team access statistics used for calculations are below:

Table 23. Red Team Access Location Basic Statistics for CIA and Aggregate Score

Paired T-Test (N = 6) for Red Team Access Location versus CIA								
	Conf. Mean	Conf. St. Dev.	Integrity Mean	Integrity St. Dev.	Avail. Mean	Avail. St. Dev.	Aggregate Mean	Aggregate St. Dev
Corporate	1.6667	1.7951	2.5000	2.5000	1.3333	1.4907	5.5000	5.5000
Operation	1.000	1.5275	2.5000	2.0616	1.000	1.5275	4.5000	4.7170

Conf = Confidentiality

Avail = Availability

Std Dev = Standard Deviation

We proceeded to test the hypothesis detailed in Section 3.1 Hypotheses:

H₀: The Aggregate scores from the Red Team access in the Corporate network are equivalent to the corresponding scores observed in the Operation network.

H_{Alternate}: The Aggregate scores are significantly different between the Red Team access in the Corporate network and the Operation network.

The significance level for this test is chosen to be $\alpha = 0.05$, where the p-values below this rejected the null hypothesis (i.e., the differentiation in scores was between access locations was significant)

A Paired T-Test was conducted to compare the overall Aggregate Score of the Red Team access in the Corporate network and the Red Team access in the Operations network. The Red Team Paired T-Test result for the Aggregate score is detailed below in Table 24:

Table 24. Red Team Access Location Aggregate Paired T-Test

Aggregate Paired T-Test	DF	Test Statistic	p-value
Sample Data	5	0.597614304667	0.576131725718

There was not a statistically significant difference in the Aggregate scores for the Red Team access in the Corporate network (M =5.5, SD =5.5) and the Red Team access in the Operations network (M =4.5, SD =4.7169906); $t(5) = 0.597614$, $p = 0.576132$.

After analysis of the aggregate scores addressing the Red Team's locations, our results suggest that the red team's location did not have significant impact on the Aggregate score. This means that there is a high probability that regardless of the location of the red team, their ability to affect a data exchange was not related to their location. Their inability, or ability, to affect the data exchanges could be accomplished equally well in either location.

6.2 Phase 2 - Operational Validation (Power Plant) Data Analysis

6.2.1 Phase 2, Operational Validation (Power Plant), Performance Objectives Descriptions

Table 25. Performance Objective 1

Performance Objective	Metric	Data Requirements	Success Criteria
Compatibility	Any plant component's failure	Position / condition of PLCs and RTUs	Continued component operation
	Any plant component's performance degradation	Position / condition of PLCs and RTUs	Continued component operation within specified tolerances

6.2.2 Performance objective 1 – Technology Compatibility

We did not use any analytical methodology in the power plant. The HECO plant engineers and NAVFAC SMEs monitored and assessed the control systems and field devices for damage, latency or other problems, and assessed results as they observed and delivered verdicts.

MANA™ and Spire were successes:

- They did NOT induce latency or performance degradation.
- They did NOT induce equipment or device failure.
- All plant systems continue to operate as they were prior to technology installation or operation.

6.2.3 Performance Objective 2 – NIDS Performance

Table 26. Performance Objective 2

Performance Objective	Metric	Data Requirements	Success Criteria
Quantitative Performance Objectives			
NIDS Performance	Pass /Fail	Caught Attack	Correctly identified >90% of attacks
		Missed Attack	Did not miss >10% of attacks

A. Performance Objective 2 – NIDS Performance:

Our analytical methodology was to compare MANA™ results with red team logs after each run. We determined on the spot if the MANA™ NIDS had caught or missed each attack.

MANA™:

- Correctly identified all (greater than the 90% criteria) of red team attacks.
- Missed no attacks exceeding the criteria of failure to identify red team activity less than 10% of the time.

6.2.4 Data Analysis

All performance was assessed as pass or fail and was visually observed by Subject Matter Experts (SME) and HECO plant engineers to determine the success or failure of the CEH / red team's efforts. Success by the CEH meant failure of the NIST or cyber defense-in-depth technologies.

The MANA™ sensor and the Spire suite were both assessed as successes.

- The MANA™ sensor was able to detect the CEH's initial reconnaissance and all subsequent attacks, with no false positives on high confidence alerts.
- The Spire technologies prevented the CEH from:
 - Compromising any replicated device within the Spire defended section of the network.
 - Giving unsafe/illegal commands to any replicated device within the Spire defended section of the network.
 - Isolating the defended devices from the control system or from performing denial of service (DoS) attacks on the ports that normally interfaced with the protected devices.

Additionally, Spire did NOT induce latency or errors in control system's communications or devices or in the plant equipment that the HMI controlled.

6.2.5 Performance Assessment

6.2.5.1 MANA™

The MANA™ NIDS was able to detect the CEH's initial reconnaissance and all follow-on attacks that the CEH later conducted. The level of success was determined by comparing the high confidence correlated alerts (NIDS output) against the CEH's attack logs to provide ground truth that there was an attack, the type of attack, when it was launched and against what IP address to determine if the NIDS provided an alert on the attack.

The MANA's patented correlation processes eliminated false positive from the "high confidence" alerts forwarded to the human operator / plant defender for action. High confidence alerts had ZERO false positives. However, lower confidence level alerts did include some false positives. The MANA™ graphical user interface was designed with the plant engineers in mind and clearly displayed the high confidence alerts for the engineers to focus on. Correlation by the multiple simultaneously running diverse algorithms reduced false positives and provided the level of confidence assigned the alert. High confidence alerts had ZERO false positives. Lower confidence level alerts and individual algorithms did include false positives. MANA's graphical user interface was designed with the plant engineers in mind and clearly shows the high confidence alerts for the engineers to focus on.

The MANA™ sensor is passive and receives all traffic via a one-way network tap. The MANA™ sensor is compatible with all networks using IP based traffic to include variations such as MODBUS over TCP/IP.

6.2.5.2 Spire Results

The CEH was unable to compromise any replicated device within the Spire defended section of the network. Commands given to devices to exceed their safety parameters were ignored by the multicomplied replications. The CEH was also unable to isolate the defended devices from the control system or to perform DoS attacks on the ports that interfaced with the protected devices. This is consistent with results from the Phase 1 designed experiment at PNNL.

Additionally, data gathered by HECO showed that Spire did NOT induce latency or errors in control system's communications or devices. In fact, plant engineers observed that in some areas Spire's response time was faster (approximately twice as fast) than the baseline system.

7.0 COST ASSESSMENT

7.1 COST MODEL

Costs for acquiring, installing and operating a cyber aware, fault and intrusion tolerant defense are outlined in the table below. Since the entire defense can be installed by the user / purchaser of the defense equipment, many of the costs in the table vary based on the labor rates of the purchaser.

Table 27. Estimated Cost Model for the Defense-in-Depth

Cost Element	Data Tracked During the Demonstrations	Estimated Costs
Hardware capital costs	Estimates made based on component costs for demonstration	\$85,000.00. (Total = MANA™ + Spire)
Installation costs (not including travel)	Labor and material required to install	- Labor costs vary depending upon who installs the defense. - Material costs = \$40 (assuming rack space available)
Consumables	Estimates based on rate of consumable use during the field demonstrations	\$600.00
Facility Operational Costs	None. (Defense does not add measurable cost to a facility.)	N/A.
Maintenance	- Frequency of required maintenance - Labor and material per maintenance action	Varies by plant location based on frequency of changes to the network or plant equipment, as well as operator preferences.
Upgrades	Estimate based on components degradation during demonstration	- Software – annually (minimum) - Hardware will not need an “upgrade” but should be replaced every three years.
Hardware lifetime	None. Estimate from SMEs.	3 years.
Operator training	N/A. Technology was operated by the technology providers.	MANA™ NIDS – Estimated eight hours. Spire System – Estimated 24 hours.
Salvage Value	N/A. No technology required replacement during the two year project.	There is no salvage value to the hardware unless the hardware is broken apart and the individual components are melted down for their precious metal content.

Hardware capital costs are the costs to acquire the server that houses the MANA™ NIDS, the server for the SIEM, the six nodes for the Spire components and four operator laptops to view / interact with the other devices and two monitors (1 x MANA™, 1 x Spire) for the plant control room.

NOTE: Users that already have a SIEM in their OT network can use that SIEM, but will still want the additional server to support the processing and memory requirements of the MANA™ NIDS. The Spire components can also be run from existing hardware but will perform better in dedicated node hardware.

If no rack space is available for the cyber defense-in-depth hardware, then there will be additional costs for racks, mounting rails, universal power supplies (UPS) and incidental consumables.

NOTE: Further experimentation is required to determine if the hardware listed above can defend one entire power plant or is best suited for only one power producing system (boiler and generator) within a power plant.

Lowest installation costs are for the user to do all work. Costs are then reduced user's labor costs to configure the Spire software for the network, install the software on the nodes and in the operator laptops, install the hardware, and tune the Spire system to the network. Installation costs for MANA™ are also only the user's labor costs to install the software, train the NIDS for the environment, install the hardware, and tune the system to the network.

NOTE: There are pre-installation costs in acquiring the network traffic, converting it into training data, cleaning the training data and training the MANA™ NIDS.

Assuming Spread Concepts installs and configures Spire, costs are estimated to be \$8,000.00 for two persons for one week, not including travel.

Assuming Resurgo installs and configures MANA™, costs are estimated to be \$15,000.00 for three persons for one week, not including travel. NOTE: Traffic to train MANA™ must be acquired AFTER installation of Spire and any other planned/scheduled network changes.

Consumables are the cables and CDs/DVDs/thumb drives/external hard drives that the user uses to download the software onto prior to loading the software into the hardware components of the defense.

Facility Operational Costs will remain essentially the same after installation as prior to installation. If all new hardware is purchased and placed into the power plant, the increase in power use and cooling requirements will be negligible assuming rack space is available.

Maintenance. Labor and material per maintenance or update action.

1. Spire: Updating/upgrading Spire requires new software to be downloaded and installed into the host devices. Spread Concepts will sell licenses and Spire support but they do not have an estimate of those costs at this time.
2. MANA™: As a rule, the MANA™ NIDS must be retrained whenever new hardware or software is added to the network. Failure to retrain will result in a significant increase in false positives and in an increased workload for the network defenders who sort thru alerts, prioritize them, and respond to them. Additionally, as attack methods and attack vectors change, MANA™ will need to be retrained to be able to differentiate the new threats from

background network traffic. Users will be able to retrain MANA™ themselves using its organic software, however, Resurgo will contract to train the NIDS or the operators and assist in the retraining. Costs are dependent upon travel and time on site at the power plant.

Electricity. Power for the new devices and monitors.

Upgrades.

- Spire software has a planned annual update/upgrade cycle as its software continues to evolve and improve at Johns Hopkins University. Spire hardware will need to be upgraded every three years in a plant environment.
- MANA™ software will not require upgrade unless significantly better algorithms are composed and released to be incorporated into the MANA™. MANA™ stays upgraded by retraining to meet new threats. MANA™ hardware should be upgraded approximately every three years in a plant environment.

Hardware lifetime. The Spire and MANA™ hardware has a normal life of five years + in the IT world. However, our SMEs suggest that within a normal power plant three years is an optimistic estimate of hardware life. Less if the hardware is placed in non-environmentally controlled areas and is exposed to high heat or moisture, degrading performance and precipitating failure. Timeframe to failure depends on the high temperature the hardware must operate in but will be less than two years. NOTE: Normal DoD information technology hardware replacement is on a three-year cycle under the current budget.

Operator training:

- Operator training will take less than twenty-four hours for Spire.
- Operator training will take less than eight hours for MANA™ installation, configuration, operation, and how to identify performance degradation, operator training can be done in eight hours. However, if the purchaser must be shown how to capture network traffic, clean the captured traffic, produce training data from the captured traffic, and train on MANA™ installation, configuration, operation, and how to identify performance degradation, then MANA™ training will take 40 hours.

Salvage Value. There is no salvage value to the hardware unless the components are broken down for their precious metals and scarce minerals. Conditions within a normal power plant will expose the hardware to adverse environmental conditions, even within environmentally controlled areas. Additionally, the cyber security of ICS systems will require the software to be shredded to ensure no IP addresses or other critical information is accidentally released from the plant.

7.2 COST DRIVERS

Technology costs are driven by processor, memory, and hard drive size. Neither MANA™ nor Spire require top of the line/cutting edge processors but quality and speed should be considered when selecting the hardware to support implementation. There are no site-specific characteristics or regional issues that will significantly impact cost.

7.3 COST ANALYSIS AND COMPARISON

Basic Site Description Assumed for The Cost Analysis

It is assumed that the technologies will be installed in environmentally controlled areas on uninterrupted power supplies (UPS) within the plants they are protecting. Cooling is important to hardware life and performance.

Assumptions Made for The Basis of The Cost Analysis

1. Hardware similar to that used in the field demonstrations will be purchased by any organization implementing the cyber defense-in-depth keeping costs similar.
2. Spread Concepts and Resurgo will continue their business model of open source software, for cost support.

Life-Cycle Cost Estimate for The Technologies

Estimating the life-cycle cost for the defense-in-depth is complicated. Information technology capacity doubles roughly every 18 months (Moore's Law). The DoD replaces its IT hardware on a three-year cycle currently. Most civilian firms replace their hardware on a three to five-year cycle to maintain par with competitors and to be able to integrate or use new software and technologies. However, historically, operational technology is used until it is no longer repairable. This includes the software and hardware required to run or interface with the antique device that has become both vital and a single point of failure. Given this, the Spire and MANATM hardware has a normal life of five years (+) when used in controlled environments. However, our SMEs suggested that conditions within a normal power plant will limit hardware life to approximately three years. Less if the hardware is not placed in environmentally controlled areas or is exposed to high heat or moisture, which will degrade performance and precipitate failure. Current cost estimate for new/replacement MANATM and Spire hardware is \$85,000.00.

The Spire and MANATM software life varies but is estimated to be six months to a year.

1. Spire software is always under review and improvement. It will be updated, and the updated / new software will be posted on the <http://www.dsn.jhu.edu/spire> website.
2. MANATM software must be retrained whenever a change is made to the network environment: new hardware, new software, new routing schemes, etc. Additionally, as attack methods and attack vectors change, MANATM will need to be retrained to be able to differentiate the new threats from background network traffic.
3. Building Life-Cycle Cost Program: The BLCC 5-3-17 does not apply to this cybersecurity project. (The NIST Building Life Cycle Cost program is available on the DOE website: <https://energy.gov/eere/femp/building-life-cycle-cost-programs>).
4. NIST Handbook 135: The Life-Cycle Costing Manual for the Federal Energy Management Program also does not apply. This cybersecurity project does not relate to the FEMP Investment Decision criteria:

- (1) It cannot be shown to be cost effective against the accept / reject system option.
- (2) It is not the most effective option for optimal efficiency because it does not impact energy or water efficiency.
- (3) It is not the most cost effective system type as it is a unique prototype and costs are greater than current best practices.
- (4) While it is an interdependent building system, it does not affect energy performance or cost.
- (5) It also does not deal with prioritization of projects. Control systems are essentially undefended currently.

Cost Comparison to Existing Technology

Assuming that DoD and commercial energy providers are extending their enterprise IT defenses over into their OT networks, costs vary depending upon organizationally negotiated prices with McAfee, Norton, Symantec or other providers for network and endpoint protection.

Costs should not vary for different regions of the country but may significantly vary based on organizational negotiating skills and organization size. The cyber defense-in-depth technology does not have to replace any existing approach but can instead be integrated with the outputs of the existing technologies using a SIEM. Integration does not save costs but use of heterogeneous sensor methodologies has been proven to increase cyber security.

Page Intentionally Left Blank

8.0 IMPLEMENTATION ISSUES

The primary implementation issue is inertia and resistance to change. Change in cyber security normally requires new tools/devices and retraining of operators to use and support the new tools. These technologies use devices that IT professionals are familiar with but they will be employed in power plants unfamiliar to IT professionals or will be operated by plant engineers unfamiliar with the equipment. For an information technology professional, implementation of the Critical Energy Infrastructure Cyber Defense-in-Depth will be no harder than any standard day-to-day IT task using tools and technology they are familiar with. Except, that they will be installing the technology in a power plant where they may have little experience or operational context. Conversely, for the plant engineer or operations technology network professional, this will be all new, even though it has been done routinely in the enterprise part of the utility's network for decades. Changes and new technology normally provoke skepticism and a reluctance to adopt the new technologies, but it can certainly be overcome by effective coordination and leadership.

The MANA™ IDS can be installed as a software package on a server in use within the power plant or for optimum performance can be installed on a new server of any manufacture. A DELL server is depicted below as an example of the type of server used during this project.



Figure 15. DELL PowerEdge R720 server

Information on MANA™ NIDS is available at: <http://www.resurgo.net>.

The Spire technologies can be installed as a software package on a server in use within the power plant. But for optimum performance, Spire should be installed on separate small nodes within the OT network. Below is a picture of Spire installed in small nodes in a rack with a standard network switch during the power plant operational demonstration.



Figure 16. Spire and Switch in a Rack In the Power Plant (4U high)

Information on Spire is available at: <http://www.dsn.jhu.edu/spire/>.

This site also contains links to additional information about Spines, Prime and other related software.

During the two demonstrations, the PNNL laboratory risk reduction test and the power plant operational validation, we learned that the MANATM and Spire technologies:

- Function in a power plant.
- Are compatible with power plant equipment.
- Do not induce latency or errors in plant networks or devices.
- Are highly effective at identifying cyber penetrations and preventing their disruptive impact.

Additionally, data gathered by plant engineers showed that Spire not only did NOT induce latency or errors in control system's communications or devices, in some areas Spire's response time approximately twice as fast as the baseline system.

REGULATIONS AND PERMITS

There are no permits required to use MANATM or Spire. However, there are regulations and codes that apply to the use of the technologies by commercial or DoD energy providers:

- 18_CFR_388-113_CEIF_2016-Vol_1, Accessing Critical Energy Infrastructure Information, 1 April 2016.
- National Institute of Standards and Technology Special Publication 800-53A, “Guide for Assessing the Security Controls in Federal Information Systems and Organizations: Building Effective Security Assessment Plans,” June 2010, as amended
- NIST Special Publication 800-82, Revision 2, Guide to Industrial Control Systems (ICS) Security, May 2015.
- Title II of the Homeland Security Act of 2002, as amended (Pub. Law 107-296): Subtitle B of title II (Authorizes DHS, through NPPD, to enhance the security, resilience, and reliability of the Nation’s cyber and communications infrastructure).
- Unified Facilities Criteria (UFC) 4-010-06, Cybersecurity Of Facility-Related Control Systems, 19 September 2016, Change 1, 18 January 2017

Regulations and codes that may apply to the use of the technologies by DoD energy providers (only):

- DoD Handbook for Self-Assessing Security Vulnerabilities & Risks of Industrial Control Systems on DoD Installations, 19 December 2012
- DoD Instruction 8500.01, “Cybersecurity,” March 14, 2014, Department of Defense
- DoD Instruction 8510.01, “Risk Management Framework (RMF) for DoD Information Technology (IT)”, March 12, 2014
- Unified Facilities Criteria (UFC) 4-010-06, Cybersecurity Of Facility-Related Control Systems, 19 September 2016, Change 1, 18 January 2017

End-user concerns, reservations, and decision-making factors.

Availability of energy production is utmost in an energy provider’s mind and is their foremost responsibility. The operational validation in a functioning power plant addressed energy industry concerns about the compatibility of, and latency induced by, the technologies.

Procurement

All hardware required is standard commercial off-the-shelf (COTS). The Spire software is open source and is built of modified open source software. MANATM is customized open source software.

Page Intentionally Left Blank

9.0 REFERENCES

1995. *Life-Cycle Costing Manual for the Federal Energy Management Program*. NIST Handbook 135.
2002. Title II of the Homeland Security Act of 2002, as amended (Pub. Law 107-296): Subtitle B of title II.
- June 2010. *Guide for Assessing the Security Controls in Federal Information Systems and Organizations: Building Effective Security Assessment Plans*. National Institute of Standards and Technology Special Publication 800-53A.
- 19 December 2012. *Handbook for Self-Assessing Security Vulnerabilities & Risks of Industrial Control Systems on DoD Installations*. DoD.
- April 2013 (includes updates as of 01-22-2015). *Security and Privacy Controls for Federal Information Systems and Organizations*. NIST Special Publication 800-53, Revision 4, Joint Task Force Transformation Initiative, Appendix B, page B6.
- 12 March 2014. *Risk Management Framework (RMF) for DoD Information Technology (IT)*. DoD Instruction 8510.01.
- 14 March 2014. *Cybersecurity*. DoD Instruction 8500.01. Department of Defense.
- May 2015. *Guide to Industrial Control Systems (ICS) Security*. NIST Special Publication 800-82, Revision 2.
- 1 April 2016. *Accessing Critical Energy Infrastructure Information*. 18_CFR_388-113_CEIF_2016-Vol_1.
- September 2016. *Cybersecurity of Facility-Related Control Systems*. Unified Facilities Criteria 4-010-06 19.
- 18 January 2017. *Cybersecurity Of Facility-Related Control Systems*, 19 September 2016, Change 1, 18 January 2017.
- 2 August 2017. *Cyber Security - System Security Management*. CIP-007-6.
2018. *Building Life-Cycle Cost Program*. BLCC 5.3-2018 (NIST Building Life Cycle Cost program. Available on the DOE website: <https://energy.gov/eere/femp/building-life-cycle-cost-programs>).

Below are websites that provide additional information related to the technologies in the defense-in-depth:

Spire: <http://www.dsn.jhu.edu/spire/>.

Spines: <http://www.spines.org/>. Spines is designed to work even when the network is compromised - that is published here: http://www.dsn.jhu.edu/papers/icdcs2016_PITN.pdf.

Prime: <http://www.dsn.jhu.edu/byzrep/prime.html>. Prime is designed to work even when the system is compromised. Published at:
http://www.cnds.jhu.edu/pub/papers/TDSC_Prime_2011_07.pdf.

In the context of SCADA, it is published here:
http://www.cnds.jhu.edu/pub/papers/Survivable_SCADA.pdf.

APPENDIX A POINTS OF CONTACT

Point of Contact Name	Organization	Phone Email	Role in Project
Kevin Jordan	Resurgo, LLC	Kevin.b.jordan@resurgo.net 808-425-5324	PI
Eamon Jordan	Resurgo, LLC	Eamon.jordan@resurgo.net 808-784-0562	Lead Test Designer
Ya'ir Amir	Johns Hopkins University	yairamir@cs.jhu.edu	JHU Task Lead (Spire, Spines, and Prime)
	Imprimis, Inc.		Imprimis Lead
Jim Brown	PNNL	jim.brown@pnnl.gov 509-375-3626	PNNL Task Lead
Luke Richards	NAVFAC EXWC	luke.richards@navy.mil 805-982-1250	Observer/Engineer
Brian Tepper	HECO	bryan.tepper@hawaiianelectric.com 808-543-7131	HECO Information Assurance Manager / Task Lead
Keith Webster	HECO	keith.webster@hawaiianelectric.com	Observer/Lead HECO Engineer
Kellie Johnsen	PACOM	Kellie.johnsen.civ@pacom.mil 808-477-8191	CTP Government representative
William J. Beary	NORAD-NORTHCOM	william.j.beary.civ@mail.mil	Observer

Page Intentionally Left Blank

APPENDIX B CALIBRATION OF EQUIPMENT

Calibration of equipment. No equipment required calibration. Sensors were “tuned” to the networks to minimize false positives, but the engineers did not consider this calibration.

Page Intentionally Left Blank

APPENDIX C QUALITY ASSURANCE SAMPLING

Quality assurance sampling. No quality assurance sampling was done during either demonstration. All PNNL network activity was recorded for later analysis back in Oahu. All HECO network activity was retained by HECO to comply with federal regulations.