

Fossil Power Plant Cyber Security Life-Cycle Risk Reduction

A Practical Framework for Implementation

2020 TECHNICAL REPORT

Fossil Power Plant Cyber Security Life-Cycle Risk Reduction

A Practical Framework for Implementation

3002019700

Final Report, November 2020

EPRI Project Manager
J. Lawrence

DISCLAIMER OF WARRANTIES AND LIMITATION OF LIABILITIES

THIS DOCUMENT WAS PREPARED BY THE ORGANIZATIONS NAMED BELOW AS AN ACCOUNT OF WORK SPONSORED OR COSPONSORED BY THE ELECTRIC POWER RESEARCH INSTITUTE, INC. (EPRI). NEITHER EPRI, ANY MEMBER OF EPRI, ANY COSPONSOR, THE ORGANIZATIONS BELOW, NOR ANY PERSON ACTING ON BEHALF OF ANY OF THEM:

(A) MAKES ANY WARRANTY OR REPRESENTATION WHATSOEVER, EXPRESS OR IMPLIED, (I) WITH RESPECT TO THE USE OF ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS DOCUMENT, INCLUDING MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, OR (II) THAT SUCH USE DOES NOT INFRINGE ON OR INTERFERE WITH PRIVATELY OWNED RIGHTS, INCLUDING ANY PARTY'S INTELLECTUAL PROPERTY, OR (III) THAT THIS DOCUMENT IS SUITABLE TO ANY PARTICULAR USER'S CIRCUMSTANCE; OR

(B) ASSUMES RESPONSIBILITY FOR ANY DAMAGES OR OTHER LIABILITY WHATSOEVER (INCLUDING ANY CONSEQUENTIAL DAMAGES, EVEN IF EPRI OR ANY EPRI REPRESENTATIVE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES) RESULTING FROM YOUR SELECTION OR USE OF THIS DOCUMENT OR ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS DOCUMENT.

REFERENCE HEREIN TO ANY SPECIFIC COMMERCIAL PRODUCT, PROCESS, OR SERVICE BY ITS TRADE NAME, TRADEMARK, MANUFACTURER, OR OTHERWISE, DOES NOT NECESSARILY CONSTITUTE OR IMPLY ITS ENDORSEMENT, RECOMMENDATION, OR FAVORING BY EPRI.

THE FOLLOWING ORGANIZATIONS, UNDER CONTRACT TO EPRI, PREPARED THIS REPORT:

Idaho National Laboratory

Southern Engineering Services

NOTE

For further information about EPRI, call the EPRI Customer Assistance Center at 800.313.3774 or e-mail askepri@epri.com.

Electric Power Research Institute, EPRI, and TOGETHER...SHAPING THE FUTURE OF ELECTRICITY are registered service marks of the Electric Power Research Institute, Inc.

ACKNOWLEDGMENTS

The following organizations, under contract to the Electric Power Research Institute (EPRI), prepared this report:

Idaho National Laboratory
1955 N. Fremont Avenue
Idaho Falls, ID 83415

Principal Investigators
S. Freeman
M. Reif
C. Rieger

Southern Engineering Services
232 Bowman Road
Aberdeen, NC 28315

Principal Investigators
B. Geddes
B. Geddes, PE

This report describes research sponsored by EPRI.

EPRI would like to acknowledge the contributions to the research and this report from the following generation member utilities:

Suzanne Bowman, LG&E and KU Energy

Brendan Buck, Oglethorpe Power Corporation

Elizabeth Holland, Duke Energy Corporation

Julie Holstad, Southern Company

Mark Jones, Prairie State Generating Company

Michael Kruger, Exelon Corporation

Michael Muzik, Southern Company

Mark Stein, Tri-State Generation & Transmission Association

Joey Ward, Duke Energy Corporation

This material is based upon work supported by the Department of Energy under Award Number(s) DE-FE0031643, and as such the U.S. Government retains a paid-up, nonexclusive, irrevocable, world-wide license to reproduce, prepare derivative works, distribute copies to the public, and display publicly, by or on behalf of the Government, this work in whole or in part, or otherwise use the work for Federal purposes.

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

This publication is a corporate document that should be cited in the literature in the following manner:

Fossil Power Plant Cyber Security Life-Cycle Risk Reduction: A Practical Framework for Implementation. EPRI, Palo Alto, CA: 2020. 3002019700.

ABSTRACT

Market conditions are forcing fossil electricity generation facility owners and operators to implement advanced digital technologies. These technologies enable efficiencies, operational flexibility, operations and maintenance efficiencies, and adapting to a transitioning workforce. These digital technologies, however, can increase the cyber security attack surface. The purpose of this research was to develop a holistic cyber security risk reduction framework for fossil generation facilities.

The framework begins with assessing how cyber risk changes across facility life cycles, including plant, system, vendor, and business life cycles. The next phase performs consequence analysis to prioritize high consequence events. Focusing on high consequence events allows owners to use a graded, risk-informed approach to prioritize cyber security efforts. The final phase identifies the digital asset attack surface in sensors and instrumentation and control equipment.

After the vulnerabilities are identified, the owner selects mitigating cyber security control measures (or countermeasures) based on the risk analysis from the previous phases. This report describes current industry cyber security best practices in fossil generation that are based on first principles for cyber security engineering. The report is divided into five sections that describe the implementation of the risk reduction framework and present identified research, methodological, and technology gaps that were identified through this course of research and development.

Keywords

Attack surface analysis

Cyber security

Cyber security first principles

Cyber security risk reduction

Operational technology (OT)

Deliverable Number: 3002019700

Product Type: Technical Report

Product Title: Fossil Power Plant Cyber Security Life-Cycle Risk Reduction: A Practical Framework for Implementation

PRIMARY AUDIENCE: Energy sector cyber security professionals

SECONDARY AUDIENCE: Technology providers and energy sector stakeholders

KEY RESEARCH QUESTION

Given the rapid change in digital controls technology and cyber attacker capabilities, how can fossil asset owners prioritize cyber security efforts to reduce risk to an acceptable level?

RESEARCH OVERVIEW

The purpose of this research was to develop a holistic cyber security risk reduction framework for fossil generation facilities. The framework begins with assessing how cyber risk changes across facility life cycles, including plant, system, vendor, and business life cycles. The next phase performs a consequence analysis to prioritize high consequence events, allowing a fossil asset owner to use a graded, risk informed approach to implement cyber security. The final phase identifies attack surface vulnerabilities in sensors and instrumentation and control equipment, also known as *operational technology* (OT), and guides the fossil asset owner to mitigate the attack surface using cyber security control mitigating measures based on the risk analysis from the previous phases. This report describes current industry cyber security best practices in fossil generation that are based on first principles for cyber security engineering. The report is divided into five sections that describe the implementation of the risk reduction framework and present identified research, methodological, and technology gaps that were identified through this course of research and development.

KEY FINDINGS

- Cyber risk changes over time.
- Cyber risk varies between the plant, system, vendor, and business life cycles.
- Consequence analysis allows a fossil asset owner to prioritize cyber mitigation efforts using a graded approach that reduces overall cyber risk at a lower cost.
- Cyber security engineering best practices are based on a set of first principles.
- Cyber security engineering best practices include consequence analysis and prioritization, identification of the attack surface, and mitigation of the attack surface to an acceptable risk level.
- There are industry challenges and research gaps that should be addressed in future research that include additional, worked examples for consequence analysis, asset and configuration management, rapidly changing regulatory and threat landscape, adapting to the rapid pace of change in the industrial control system and cyber security technologies, event response, secure configuration for OT assets, and supply chain integrity and testing.

WHY THIS MATTERS

Fossil generation assets are targets for sophisticated nation state and other cyber adversaries. A cost-efficient cyber security framework is needed to mitigate cyber risk to an acceptable level that fits within the business constraints of fossil asset owners.

HOW TO APPLY RESULTS

Develop a cyber program based on the cyber security framework and best practices using the three-phased approach as described in the report.

LEARNING AND ENGAGEMENT OPPORTUNITIES

- Collaborate with the Electric Power Research Institute (EPRI) Cyber Security Program.
- Attend an EPRI Technical Assessment Methodology (TAM) training session.
- Attend an EPRI Cyber Security in Supply Chain, Cyber Security Procurement Methodology training session.
- This product is useful to fossil asset owners, equipment and technical service vendors, industry standards bodies, and government agencies associated with critical infrastructure cyber security.

EPRI CONTACT: Jeremy Lawrence, Senior Technical Leader, jlawrence@epri.com

PROGRAM: Cyber Security for Generation Assets, P209

Together...Shaping the Future of Electricity®

Electric Power Research Institute

3420 Hillview Avenue, Palo Alto, California 94304-1338 • PO Box 10412, Palo Alto, California 94303-0813 USA

[800.313.3774](tel:800.313.3774) • [650.855.2121](tel:650.855.2121) • askepri@epri.com • www.epri.com

© 2020 Electric Power Research Institute (EPRI), Inc. All rights reserved. Electric Power Research Institute, EPRI, and TOGETHER...SHAPING THE FUTURE OF ELECTRICITY are registered service marks of the Electric Power Research Institute, Inc.

CONTENTS

ABSTRACT	V
EXECUTIVE SUMMARY	VII
1 BACKGROUND AND INTRODUCTION	1-1
Background	1-1
Introduction	1-2
Purpose.....	1-5
How to Use the Framework	1-6
2 LIFE-CYCLE FRAMEWORK FOR CYBER SECURITY RISK REDUCTION	2-1
Temporal Nature of Life Cycles	2-1
Plant Life-Cycle Temporal Qualities	2-2
Business Owner Life-Cycle Temporal Qualities	2-3
System and Component Life-Cycle Temporal Qualities	2-3
Vendor Business Temporal Qualities	2-4
Cyber Security Life-Cycle Temporal Qualities	2-5
Frequency.....	2-5
Life-Cycle Influences on Cyber Security	2-6
3 CONSEQUENCE ANALYSIS AND PRIORITIZATION: CRITICAL FUNCTIONS AND SERVICES THROUGH LIFE CYCLES	3-1
CCE Consequence Analysis Best Practices	3-2
TAM Consequence Analysis Best Practices	3-5
Other Consequence Analysis Methods.....	3-6
EPRI Guideline for Establishing Criticality Rankings for Fossil Power Plant Systems and Components (1025341).....	3-6
IEC 62443 Security for Industrial Automation and Control Systems	3-7
IEC 61508/61511: Functional Safety of Electronic Systems.....	3-7

Consequence-Based Critical Function Prioritization Using CCE	3-8
Introduction to Hypothetical HCEs	3-9
HCE #1—Lubrication Oil System Failure.....	3-10
HCE #2—Communications System Disruption.....	3-12
4 BEST PRACTICE STRATEGIES FOR CYBER ENGINEERING	4-1
Cyber Security First Principles.....	4-2
Systems Engineering First Principles of OT Cyber Security	4-3
Risk Management First Principles for OT Cyber Security.....	4-4
Cyber Security Engineering and Design First Principles	4-4
CCE and TAM as First Principle Best Practices.....	4-5
CCE Overview	4-5
EPRI TAM Overview	4-7
Bounding the Universe of Software and Hardware Vulnerabilities.....	4-9
Mitigation of the Entire Kill Chain	4-9
TAM Implementation Use Cases	4-9
The Complementary Relationship of CCE and TAM	4-10
Best Practices for Identifying Attack Surface Vulnerabilities	4-10
CCE Attack Surface Identification Best Practices.....	4-12
TAM Attack Surface Identification Best Practices.....	4-15
IDART Attack Surface Identification Best Practices	4-17
Best Practices for Risk-Informed Attack Surface Mitigation	4-18
CCE Risk-Informed Attack Surface Mitigation Best Practices	4-20
TAM Risk-Informed Attack Surface Mitigation Best Practices	4-23
Control Method Security Effectiveness.....	4-23
Combined Security Effectiveness	4-23
Risk-Informed Target Level Threshold	4-24
Control Method Efficacy	4-24
Allocate Preferred Engineered Security Control Methods.....	4-25
Shared Site-Level Security Control Methods and Relationship Sets.....	4-25
IDART Risk-Informed Attack Surface Mitigation Best Practices	4-26
CCE Best Practice Case Study	4-26
Baltavia Power Outage [30]	4-27
TAM Best Practice Case Studies.....	4-29
Rapid and Effective Cyber Incident Response Using the TAM.....	4-29

TAM Cyber Engineering Approach Saved Over 1 Million Dollars in a Procurement	4-30
Thoroughness and Efficiency Using the TAM Cyber Engineering Approach	4-30
5 GAP ANALYSIS AND RESEARCH RECOMMENDATIONS.....	5-1
Identifying Challenges and Research Gaps.....	5-2
Standards, Methodologies, and Guides	5-2
Tools and Technologies.....	5-3
When to Use Each.....	5-3
Challenge and Gap Analysis Criteria	5-3
Industry Challenges, Notable Gaps, and Research Needs	5-4
Applying a Graded Approach to Cyber Security	5-4
Asset and Configuration Management	5-4
Changing Regulatory Landscape	5-6
Rapidly Changing Technologies and IT/OT Convergence	5-6
Event Response.....	5-6
Secure Configurations for OT Assets	5-7
Supply Chain Integrity and Testing	5-8
6 REFERENCES, BIBLIOGRAPHY, AND ABBREVIATIONS.....	6-1
References.....	6-1
Bibliography	6-3
Abbreviations and Acronyms.....	6-3
A CYBER SECURITY TOOLKIT—EXAMPLES OF STANDARDS, GUIDES AND METHODOLOGIES, AND TOOLS AND TECHNOLOGIES.....	A-1
Standards, Guides, and Methodologies.....	A-1
Technical Assessment Methodology.....	A-1
Consequence-Driven Cyber-Informed Engineering.....	A-1
System-Theoretic Process Analysis.....	A-2
EPRI HAZCADS	A-2
EPRI Component Criticality Ranking.....	A-2
Cyber Process Hazards Analysis.....	A-2
North American Electric Reliability Corporation CIP	A-3
IEC 62443.....	A-3
ISO 27001	A-3
NIST Cyber Security Framework	A-3

NIST SP 800.....	A-3
IEC 61508 and IEC 61511	A-4
CISA Assessments and CSET	A-4
Cybersecurity Capability Maturity Model	A-4
Technology Capabilities	A-4
Hardware Integrity Checking.....	A-5
Penetration Testing.....	A-5
Intrusion, Prevention, and Anomaly Detection Tools IDS/IDPS/HIDS	A-5
Security Information and Event Management	A-6
Event Response and Automation.....	A-6
Identity and Access Management.....	A-6
Secure Remote Access	A-7
Network Boundary Protection and Network Access Control (NAC)	A-7
Asset and Configuration Management.....	A-7
Computerized Maintenance and Materials Management	A-8
Compliance Management	A-8
Encryption and Data Protection	A-8
Software and File Verification	A-8
Backup and Recovery.....	A-9

LIST OF FIGURES

Figure 1-1 High-level framework process	1-6
Figure 2-1 Notional illustration of event frequencies	2-6
Figure 2-2 Notional life cycles and their influence on cyber security	2-7
Figure 3-1 Image comparing the impact of HCE #1 on a facility during three different stages of its life cycle.....	3-11
Figure 3-2 Image comparing the impact of HCE #2 on a facility during three different stages of its life cycle.....	3-13
Figure 4-1 High-level overview of the four phases in the CCE methodology	4-6
Figure 4-2 TAM overview	4-8
Figure 4-3 The CCE kill chain describing Phase 1 through Phase 3.....	4-14
Figure 4-4 Assessment scope.....	4-15
Figure 4-5 Activities to identify the installed configuration and data flow	4-16
Figure 5-1 Gap analysis process.....	5-1

LIST OF TABLES

Table 4-1 Security control method efficacy scoring	4-25
Table 4-2 Severity scoring matrix of the Baltavian case study HCE (cyber-event #4).....	4-28

1

BACKGROUND AND INTRODUCTION

This report is the culmination and output from research funded by the U.S. Department of Energy (DOE), Office of Fossil Energy. The purpose of the research was to develop a holistic cyber security risk reduction framework for fossil generation facilities. The framework implements identifying attack surface vulnerabilities in sensors and instrumentation and control (I&C) equipment, also known as *operational technology (OT)*, guiding risk-informed cyber security control mitigating measures, and understanding current industry best practices in fossil generation. The report is divided into five sections that describe the implementation of the risk reduction framework and present identified research, methodological, and technology gaps that were identified through this course of research and development.

Background

The state of cyber security for fossil generation plants is rapidly changing. Generation sector-specific challenges and drivers exist that continually put pressure on asset owners and operators to maintain secure OT operations. However, the rapidly changing operational factors coupled with challenges and drivers make it difficult to know what equipment to secure in the plant and the best way to secure it.

Nation-state adversaries are increasingly becoming more sophisticated in their tactics, techniques, and procedures (TTPs) for conducting cyber operations against critical infrastructure. The most advanced adversaries have an expanded understanding of industrial control system (ICS) and OT equipment, communications protocols, and networking. They are able to target, perform reconnaissance, craft specialized malware, and deliver payloads to ICS and OT devices. As they continue to increase in sophistication, fossil asset owners and operators must continually secure their ICS and OT networks from new TTPs. Examples of sophisticated attacks on critical infrastructure ICS have included STUXNET, BLACKENERGY, INDUSTROYER, and TRISIS [1]. Each event had disrupting consequences on their equipment, systems, and facilities. In addition to nation-state actors, there has been an increase in commodity malware and ransomware targeting large industrial companies and utilities. In February 2020, a natural gas pipeline was impacted by ransomware that caused a section to be taken out of service and a loss of production and revenue over two days [2].

Market conditions are forcing fossil asset owners and operators to become more efficient, more flexible (cycling), reduce operations and maintenance (O&M) costs, and adjust to a transitioning workforce. Efficiency and flexibility are linked in current market conditions that require fossil plants to have a high heat rate, a fast ramp rate, lower the minimum load, and be able to cycle multiple times a day when dispatched. Enhancements in digital technologies, especially within the control system, are needed to accomplish these tasks. As more digital technologies and capabilities are added to the control systems and OT networks, they must be configured and hardened to minimize the attack surfaces that could be exploited by an attacker. New process

and equipment monitoring capabilities are being installed in plants to aid in increasing efficiency and reduce O&M annual expenditures. The data from sensors and instrumentation are being aggregated and sent to corporate monitoring and diagnostics centers for analysis. The additional monitoring capability allows engineers to have more data on equipment operating conditions and process variables. When the variables become out-of-band or operating conditions show anomalous readings, the equipment can be repaired before a failure, or routine maintenance can be delayed until conditions warrant. As these data connections are made from plant ICS and OT networks to networks outside the plant, new attack surfaces could be created that could be exploited by an adversary.

Additionally, resource availability, resource shortages, and displaced expertise continue to be problematic for many fossil generation utilities. These challenges have also been amplified by the recent COVID-19 pandemic, causing utilities to have reduced on-site staff or sequestered staff. To address these challenges, technologies like secure interactive remote access for utility personnel and vendors that require logical access to ICS and OT equipment.

ICS and OT equipment original equipment manufacturers and vendors continually develop new products, upgrades, and versions. The typical distributed control system (DCS) upgrade cycle for fossil plants is between 5 and 10 years. Other subsystems or individual equipment components are updated on a more regular basis. Many of these new technologies help utilities address their marketplace pressures. Some technologies come from the information technology (IT) domain and are implemented in OT networks, such as virtualization and security equipment. Other technologies are adopted in the fossil generation sector from other critical infrastructure ICSs. As new and emerging technologies are implemented at fossil plants by utilities and vendors, the equipment and the networks must be assessed, configured, and secured in order to prevent creating additional attack surfaces and to mitigate the attack surfaces that could be exploited by an adversary.

Last, the cyber security regulatory landscape is continuously evolving. The North American Electric Reliability (NERC) critical infrastructure protection (CIP) standards contain the cyber security regulatory requirements for fossil generation plants. There are currently 14 NERC CIP standards. Not all standards and regulatory requirements apply to all fossil plants. New standards are being drafted, and revisions to current standards are ongoing. These new and revised standards contain additional regulatory requirements for compliance. As NERC regulatory compliance requirements continue to evolve, fossil plants need to react to the new requirements continually.

All of these generation, sector-specific challenges create compounded cyber security challenges that fossil plants need to address to ensure that they can protect from a cyber attack, detect an ongoing cyber attack, and respond and recover from a cyber attack.

Introduction

Risk is the product of the consequence of a successful cyber attack and the likelihood that a successful cyber attack occurs. Risk can be qualitative or assigned numerical values and categorized quantitatively. The overall risk of a successful cyber attack can vary, based on numerous factors. Some factors could include the function or criticality of the component or system, which impacts the consequence. For example, two of the same make and model programmable logic controllers (PLCs) could have the same likelihood of a successful attack. Still, the function that they serve in the plant can be different (one controlling a critical process

and the other controlling a non-critical process), which could have different consequence levels impacting their risk profile. For this reason, it is essential to have high fidelity insight into the plant's asset management activities and understand what digital equipment is installed, their communication pathways, and their access levels. It's also essential to perform a criticality analysis of the equipment to understand its plant function and if it could impact safety, security, reliability, or a critical plant process.

In addition to varying risk for systems and equipment in the plant, risk can also be time-dependent. Depending on the life cycle of the plant, the risk of a successful cyber attack can be higher or lower. For example, if a cyber attack is successful during the construction or commissioning phases of a new plant, the risk could be lower than while the plant is operating at load because the consequence could be higher. If the plant is in an outage, the consequence of a cyber attack may be lower because it is not at load, but the likelihood may be higher because more personnel (outage crews) are physically on-site with more physical and logical access to digital plant equipment. Plants that are on a glide path to decommissioning and closure may have a lower risk because they may not be considered a critical asset for reliability, and the consequence of a cyber attack could be lessened. It is important to note that cyber risk can never be reduced to zero. There is always some sort of consequence associated with a successful cyber compromise, and the likelihood can never be eliminated (other than not having operational digital equipment).

This should not be confused with the risk acceptance or risk tolerance level. Risk acceptance is unique to each utility, system, and asset in a utility. Many factors can contribute to a risk acceptance decision, including but not limited to the following:

- Reliability requirements (critical loads)
- Regulatory requirements
- Risk exposure history
- Potential consequences (for example, environmental or public impacts)
- Financial considerations

The risk acceptance level should be set, and risk above that level should be mitigated—meaning that the likelihood of a successful cyber attack or the consequence of a successful cyberattack should be mitigated down to the acceptable level. To mitigate the cyber risk, fossil utilities determine the individual risk associated with digital systems and equipment at each asset. They determine the likelihood of a successful cyber attack and the consequence of a successful cyber attack, as mentioned previously. To fully understand the likelihood and consequence, attack surface vulnerabilities in those systems and equipment must be identified. Attack surface vulnerabilities are mitigated through the application of cyber security controls. Mitigation of vulnerabilities lowers the likelihood and therefore the risk of a successful cyber attack.

While there are many tools, standards, and methods to determine the attack surface vulnerabilities (in this context, the set of vulnerabilities that comprise the attack surface that are more than published software flaws and vulnerabilities) associated with digital OT equipment, it is critical that any approach is an implemented aspect and is implemented consistently. For many of the standards and methods, the process of attack surface vulnerability identification can be resource-intensive and require input from multiple subject matter experts (SMEs) and vendors. For these reasons, fossil plants may not have completed a thorough attack surface vulnerability assessment

on their digital systems and equipment. While basic asset and configuration management processes continue to be a challenge across the entire generation sector, it is even more challenging to complete an attack surface vulnerability assessment on those same systems and equipment. If there are not accurate asset inventories and plants don't have an accurate baseline configuration documented, it can be difficult and time-consuming to understand which vulnerabilities exist.

Security controls are engineered to mitigate attack surface vulnerabilities. In addition to understanding the systems and equipment along with the baseline configuration, fossil cyber practitioners must also understand the features, functions, and capabilities of the equipment. These features, functions, and capabilities can be used and configured, in many cases, to mitigate some identified vulnerabilities. These could be features such as configuration change flags, setting passwords, disabling ports, configuring allowable network traffic, media access control (MAC) address locking, and so on. Many of these technical security controls fall into the category of cyber hardening activities. For attack surface vulnerabilities that cannot be mitigated by technical means, they may need to be mitigated administratively, that is, through process, procedure, or best practice.

There is an abundance of tools, standards, and methodologies to determine risk levels and perform attack surface vulnerability assessments. For the allocation of security controls, there are numerous technical tools, technologies, software, and hardware devices that can provide protection, detection, response, and recovery functions. Trying to determine which tool, standard, or methodology to use in a fossil generation plant can be overwhelming for cyber practitioners. It is essential to use a process that is right-sized, efficient, complete, and repeatable. The process for applying the standards and methods can be resource-intensive because there generally are no automated data collection and analysis tools. These tasks require utility SMEs to work with vendors to collect and analyze configuration and equipment data for each system. Often, there are multiple roles at the utility that are involved with these efforts because not one single person or position has all the necessary information required for the complete analysis. This tends to be problematic because there are typically few dedicated personnel assigned solely to cyber security for OT. Many of the resources are I&C engineers who also perform secondary cyber security tasks.

These considerations present the current state of generation sector cyber security in ICS and OT systems. Several cyber security process and technology gaps have been identified that could significantly improve the efficiency and consistency of vulnerability identification and mitigation within fossil generation. Because personnel typically performing these tasks have primary job functions other than cyber security or there is not enough dedicated staff to perform these tasks, automation and tools are still needed. Automation can also provide the ability to counter cyber attacks. An intelligent and automated response reduces the response timescale by removing humans from the response loop. The hope would be that, at some point in the future, autonomous tools will be available to provide efficiency to generation OT cyber security practitioners for their roles and responsibilities.

Research and tool gaps have been identified as part of this research and development effort. Those are specifically reported in Section 5 of this report. There are already numerous methods, standards, and tools available to aid cyber security practitioners. However, there are also several identified opportunities for research and development to enhance the efficiency of staff or to provide enhanced protection, detection, response, and recovery functions.

Purpose

The objective of this research is to provide a fossil generation sector-specific framework that utilities and plants can employ to ascertain an acceptable risk level, identify attack surface vulnerabilities, and allocate cyber security controls in ICS and OT systems. The framework is employed to reduce the cyber risk of fossil plant's ICS and OT systems to an acceptable or tolerant level. The framework's goals are to be the following:

- Efficient
- Flexible
- Repeatable and consistent
- Performance-based

The framework was designed to be efficiently implemented. It is concise to minimize jargon and theory and is written using plain English. It was recognized that some fossil plants do not have dedicated OT staff or extensive experience securing ICS. For this reason, the framework helps implement cyber security through all of the plant's life cycle phases and builds upon previous cyber security controls put in place to protect, detect, respond, and recover from a successful cyber attack. It acknowledges the division of responsibilities for the different roles in the plant or utility along with vendors, suppliers, and contractors so security control allocation efforts can be maximized.

There are numerous standards, tools, and methodologies that fossil plants can choose to implement to determine risks and identify applicable attack surface vulnerabilities. This framework provides maximum flexibility by not specifying which methodology any plant should employ, and instead provides the basis to how the methodologies should be selected, which elements are necessary, and how they should be applied. It also does not create new methods or duplicate any other processes but leverages what has already been tested and applied in ICS and OT networks. For illustrative purposes, this framework used the Idaho National Laboratory's (INL) Consequence-driven Cyber-informed Engineering (CCE) Methodology [3] and the Electric Power Research Institute's (EPRI) Cyber Security Technical Assessment Methodology (TAM) [4]. These methodologies can be used to identify vulnerabilities and allocate cyber security control methods; however, other methodologies, standards, and best practices (for example, from the National Institute of Standards and Technology [NIST], International Electrotechnical Commission [IEC], regulatory frameworks, and so on) could also have been selected.

The framework is repeatable and consistent, meaning that one person at the utility should be able to repeat the process and obtain the same results, or multiple personnel using the framework yield the same results. Also, they should be able to implement the framework across an entire fleet to yield consistency. The framework itself is not complicated and presents the overall process in a way that it is able to be tailored to the actual implementation methodologies chosen.

The framework is performance-based. Because it does not dictate which standards, tools, or methods to use, it can be accomplished by performing the basic phases outlined in Sections 2–4. By performing a temporal risk analysis, vulnerability analysis, and cyber security control allocation, cyber security practitioners are able reduce the cyber risk within the plant rather than simply apply a checklist or catalog-based cyber control regimes, such as regulatory compliance.

How to Use the Framework

The framework is divided into three phases, as presented in Sections 2–4. All three phases are used in sequence to identify the acceptable risk level for the fossil plant, systems, and equipment; determine the consequences of a successful cyber attack; identify attack surface vulnerabilities in plant systems and equipment; and allocate security controls to mitigate the vulnerabilities.

Figure 1-1 illustrates the phases of the framework. Within each phase, several actions are recommended to determine risk, understand critical functions and consequences, and identify and mitigate the attack surface.




Figure 1-1
High-level framework process

Plant owners begin by determining the risk of a cyber attack associated with the life cycle phases (Section 2). This risk may be different for each life cycle phase and could vary from plant-to-plant or utility-to-utility. The risk for systems and plant equipment could also be different from the overall risk level of the facility and be temporal in nature, possibly changing several times over its life span. In this phase, management uses the risk levels to determine the acceptable risk. The risk acceptance level has numerous aspects that factor into the final decision and is used later in the framework to understand the appropriate number and strength of vulnerability mitigations that need to be implemented as part of a defense-in-depth strategy.

Building upon the life cycle risk, cyber practitioners perform a consequence analysis to determine and prioritize the critical plant functions and the equipment that is required to maintain those critical functions (Section 3). The critical plant functions vary from plant to plant, based on loads, plant type, plant manufacturer, operations considerations, and many others. For example, if a co-generation plant were applying the framework, it might determine that steam production

was a more critical function than electricity generation. If a plant serves a critical load, it may have requirements to provide a certain amount of base power. After critical plant functions and equipment are identified, the consequences of a successful cyber attack on those critical components are evaluated and the high consequence events (HCEs) are identified. The TAM and CCE processes were studied to illustrate how a cyber security practitioner can implement these approaches, but other tools, standards, and best practices could be employed to yield the same or similar results.

Section 4 discusses applying best practices for cyber security engineering. Best practices are a collection of tools, processes, or standards that utilities can employ in-part or whole to help identify the attack surface vulnerabilities present in digital plant systems and equipment. Several best practices can be used if first principles are maintained. Again, in this phase, the TAM and CCE, along with Sandia National Laboratories Information Design Assurance Red Team (IDART) methodologies, were highlighted as best practices that can be used to determine the attack surface vulnerabilities and methods to identify controls to mitigate the known vulnerabilities. Using the previous phases to apply mitigations that are risk-informed is a key consideration. Users consider the security control effectiveness and match it with its implementation burden to determine whether it should be implemented.

A defense-in-depth approach is used to apply security controls that protect, detect, respond, and recover digital assets. Appendix A lists a compendium of available tools, methodologies, and standards that can be used by fossil plants to implement a cyber security program (for example, technical control application), determine risk, or identify attack surface vulnerabilities and mitigation techniques.

The final section (Section 5) contains identified technological and research gaps that were identified throughout the research process. The gaps are presented to the framework implementer as challenges that exist in the sector and areas for future research to increase the overall cyber security posture of the industry and fossil plants. Many of the items that have been identified are ancillary to cyber security but must be implemented at a high level to accomplish high cyber effectiveness. There were also a number of technology gaps specific to fossil generation that could greatly enhance cyber security operations, cyber harden plant equipment, and make fossil generation plants a less desirable target for an attack. Information in this section could be used to develop an overall research roadmap for future cyber security development.

2

LIFE-CYCLE FRAMEWORK FOR CYBER SECURITY RISK REDUCTION

Fossil generation facilities manage and mitigate multiple risks over their lifetimes, including safety, environmental, regulatory, and cyber risks while also meeting generation performance objectives within an increasingly competitive market. This section provides a framework for reducing risks within the context of the facility life cycle.

According to Merriam-Webster, a *life cycle* is “a series of stages through which something [a fossil generation facility] passes through its lifetime.” The overall facility life cycle is decomposed in this report into more refined life cycles that shed light on specific performance and risk issues. The more refined life cycles include the following:

- Plant systems and components, including subsystems and software
- Digital product life cycles
- Cyber security life cycles, including kill-chains, security control method persistence, and threat intelligence

Cyber risk and risk management may vary between life cycles and within a life cycle. Each life cycle has unique characteristics that affect risk, and cyber risk may vary over time within a life cycle. Understanding how risk varies within a life cycle allows a plant owner to more effectively manage risk. The cyber strategy and program should reflect the influence of life cycles and may evolve with life cycles.

Sections 3 and 4 describe best practice approaches for identifying life cycle risk, cyber vulnerabilities, and cyber vulnerability mitigation using a risk-informed approach.

Temporal Nature of Life Cycles

The common denominator across all life cycles is time, meaning that all life cycles are temporal in nature. According to Merriam-Webster, the word *temporal* is an adjective that is “of or relating to the sequence of time or to a particular time.” The temporal nature of each type of life cycle is different and allows a plant owner to gain insight into how cyber risk can be managed in the context of time within a life cycle. This subsection describes the temporal qualities of key life cycles. Another temporal quality is the frequency of periodic events within a life cycle or across life cycles.

Plant Life-Cycle Temporal Qualities

Fossil generation facilities are designed for a certain lifetime. The designed lifetime and point in time of the facility life drive decision-making. At the same time, market prices, facility performance, regulatory forces, and costs may cause an owner to shorten or extend the facility design life. Shortening a plant life may allow the owner to accept higher cyber risk and lower costs.

Facility life extension is usually characterized by the following:

- System/component replacements or upgrades
- Advanced monitoring and maintenance practices, including management of age-related degradation mechanisms
- Advanced failure analysis
- Advanced operating practices

Most life extension approaches involve some level of digital technology that expands the cyber attack surface of the facility and should be considered in the cyber strategy and program.

Conversely, system and component aging, operating costs, regulatory pressures, and market pressures can lead to the early retirement of a facility.

Example 2-1 Facility near the end of life

Consider a facility that is near the end of its life cycle and has been designated for shutdown and decommissioning within the next four years. The owner has made the business risk decision that non-safety significant systems and components are allowed to run to failure (RTF), and if those systems or components fail and require significant repair or replacement to restart, the plant will be permanently shut down.

The risk tolerance for equipment failure is higher near the end of the facility life cycle, resulting in lower operating costs through the end of life.

The higher risk tolerance can also be translated to cyber security in areas such as digital system upgrades and patches. A vendor may offer new cyber security capabilities through an upgrade; however, the higher risk tolerance allows the owner to forego the purchase. The same risk decision applies to software security patches, allowing the owner to forego the installation of security patches.

Business Owner Life-Cycle Temporal Qualities

A business that owns a generation plant also has a life cycle. Businesses are affected by a variety of factors that force the business through renewal or end of life. Financial, legal, and regulatory performance are factors that influence the business life cycle. Business performance changes over time and is regularly reviewed for impacts on the business risk. Corporations and corporate boards have risk committees. Risk committees have recognized that cyber security risk needs to be considered as part of the overall business risk and are beginning to influence cyber security strategy and programs at the generation facility level. Business performance is reviewed in the context of risk, including cyber security risk, and may affect the facility life cycle or risk tolerance.

Other factors that influence the business life cycle include the business model, investor goals, market forces, regulatory forces, and acquisition or divestiture. Each of these factors vary over time and may drive changes in cyber security strategy and programs.

Example 2-2 Business model change

Consider an owner that changes its business model from owner/operator to owner-only by outsourcing facility operations. The new operator uses advanced maintenance techniques and remote access to boost operational efficiency and thus reduce O&M costs. The owner decides that additional security control methods must be implemented to reduce the risks due to the remote access expected by the new operator.

Also, the plant boundary and network must be modified in order to allow remote access. The system modifications and introduction of additional attack pathways increase the cyber security attack surface of the plant, thereby increasing risk. The board's decision to lower risk tolerance and implement business model change to allow remote access requires changes to the cyber security strategy and program to provide additional risk mitigation.

System and Component Life-Cycle Temporal Qualities

System and component life will vary with overall facility life, from system to system, and from component to component. System and component life may be dictated by the following:

- Equipment aging and performance
- Obsolescence and vendor support
- Obsolescence and availability of replacement parts
- Availability of qualified personnel to operate and maintain the system or component
- Compatibility with other system or component upgrades

Facility personnel are the most knowledgeable of system and component life cycles. As digital systems and components age, operators and maintenance personnel must manage the issues listed above. A cyber security strategy that considers system and component life cycles is more flexible, reduces costs, and provides more significant mitigation.

Example 2-3

Component obsolescence

Consider a digital component (such as a controller) that is no longer supported by the vendor. However, the controller has functioned flawlessly for 15 years, and engineering and plant staff do not believe the cost of an upgrade is justified because the maintenance history suggests that failure is not imminent, and an upgrade will not improve plant performance.

However, the controller configuration software only runs on the Windows XP Operating System (OS). A maintenance laptop running Windows XP must be maintained and operated in the facility. Because Windows XP is no longer supported and patched by Microsoft, XP is easily exploited by an adversary. Also, modern anti-malware, whitelisting, and other cyber security tools are not compatible with XP. The result is a maintenance laptop and digital component with an expanded attack surface that is not able to use preferred tools.

The facility's cyber strategy and program deals with these types of life-cycle issues by adopting specific methods that are compatible with XP, such as additional physical security access controls, physical access control logging, and XP system recovery.

Vendor Business Temporal Qualities

A vendor's business life cycle can have a significant impact on the plant owner's business, plant systems, and components. A vendor develops a business strategy that leads to a product strategy. Competition, financial performance, disruptive technologies, acquisition, technical skills (or lack thereof), recruiting and retention, and other factors drive a vendor's business life cycle. Vendor business life-cycle changes can greatly affect other life cycles, such as a vendor's product life cycle or an owner's facility or specific system's life cycle.

Product research and development, product and version support, and software patches to fix security vulnerabilities are all affected by the vendor's business life cycle. Increased reliance on digital systems and components, increased outsourcing for digital system and component support, and product life-cycle changes can have a significant impact on the facility's cyber security strategy and program.

Example 2-4

Vendor acquisition

Consider a digital system sold and supported by a vendor that has been acquired by a competitor. The competitor acquired the vendor for several business reasons and has a competing product line that is redundant to the original vendor's product installed in many generation facilities. The competitor (now the new owner) has decided to discontinue development and support of the original vendor's product line and is encouraging facility owners to switch to the competing product line.

The original vendor's product line is not obsolete and is up to date; however, its business life cycle has been disrupted, resulting in a forced end of the original product life cycle.

Cyber Security Life-Cycle Temporal Qualities

Cyber security strategies and programs also have life cycles. New software or hardware vulnerabilities are discovered or introduced, adversary skills and capabilities change, new technologies are implemented, and the skill level and resources required to overcome a security control method may change over time. These issues are compounded by a changing threat landscape measured through threat intelligence, and cyber intrusion dwell times typically ranging from 6 to 18 months.

The effectiveness of some security control methods also decreases over time. For example, the longer a password is in use, the more likely the password will be compromised. Therefore, as the exploit difficulty decreases over time, the likelihood (an element of risk) of a successful attack increases. This report uses the term *persistence* to indicate the effective duration of security control methods.

Example 2-5 Unpatched Windows Server

Consider an MS Windows Server configured as a DCS Application Server on an Active Directory (AD) domain. The DCS and AD accounts are primarily static accounts with a limited number of new and disabled account changes on an annual basis.

When last updated, the DCS Application Server was fully patched. However, due to ongoing operational constraints, the DCS Application Server is not patched again until the vendor releases a new update and the update is scheduled for installation during an outage. Windows Server vulnerabilities are discovered almost daily and Microsoft patches the Windows Server once and sometimes twice monthly. The longer that the Windows Server goes without a patch, pending an outage, the lower the skill level needed by an attacker to overcome the Windows Server control methods, because information about a vulnerability and crafted exploits become more widely available.

If operational constraints delay software updates and patches, then the cyber program should include additional security control methods, such as isolation from the business network, to mitigate system vulnerability risk that increases over time.

Frequency

Within a life cycle, certain activities or events occur with more or less frequency. The frequency of published MS Windows Server patches is high, whereas the frequency of vendors going out of business is low. When developing a strategy to mitigate potential cyber security risks associated with one or more life cycles, the frequency of certain activities or events should be considered. Figure 2-1 illustrates frequency over the time of different events.




Figure 2-1
Notional illustration of event frequencies

Life-Cycle Influences on Cyber Security

Figure 2-2 illustrates how various life cycles can “stack up” on a common timeline, in this case, a notional, baseload, coal-fired generation facility. Note that Figure 2-2 does not represent any actual facility.




Figure 2-2
Notional life cycles and their influence on cyber security

The overall time frame in Figure 2-2 is 25 years. The current year is noted as “Year N” at 5 years into the timeline. N-5 indicates 5 years ago, N+5 indicates 5 years into the future, and so on.

The overall facility life cycle is indicated at the top of Figure 2-2. The colors represent facility profitability, where green indicates profitability, yellow indicates break-even, and red indicates losses. The facility life cycle shows several types of events, including planned and forced outages, temporary curtailments, increasing or decreasing market pricing conditions, increasing or decreasing fuel costs, increasing labor cost, and, finally, end of life, marked by increasing labor costs, increasing fuel costs, and decreasing market pricing.

Plant system life cycles are indicated below the facility life cycle. Only three systems are shown for brevity, but they illustrate some key points. The colors represent system health in terms of reliability and availability, where green represents good health, yellow indicates weaknesses, and red indicates unhealthiness. Note that cyber security “health” is indicated at the bottom of Figure 2-2, discussed below. Plant system life-cycle colors in Figure 2-2 represent only reliability and availability.

The boiler life cycle shows emerging weakness in year N+3 and then unacceptable health in year N+4, resulting in a forced outage (shown on the facility life cycle) for retubing major portions of the boiler due to a significant failure event. Weakness emerges again in year N+8, followed by a planned outage in year N+9 for inspections and repairs. The cycle is repeated in years N+13 and N+14.

The turbine generator life cycle shows that inspection and repairs are sufficient during the forced outage in year N+4, but a major overhaul of the turbine is necessary in the planned outage in year N+9, and a generator rewind is necessary in year N+14 during another planned outage.

The OT life cycle shows that a system refresh was performed last year (N-1) to resolve some obsolescence issues and, today, OT system reliability and availability are healthy. A system security update, available and on standby for a system outage, is installed during a facility curtailment period (shown on the facility life cycle) about halfway through year N+1. Two more refreshes occur in years N+4 and N+9, during respective forced or planned outages. Another update is shown in year N+11, installed during another curtailment period, and, finally, a decision to maintain the OT systems without any further refreshes is made during a planned outage in year N+14. The facility owner wants to wait and see whether fuel costs and market conditions can overcome the increasing labor costs before making any more capital investments, especially following a rewind of the main generator.

The OT product life cycles, under the control of the vendor, are shown below the plant system life cycles in Figure 2-2. Three generations of the OT product are shown (X, X+1, and X+2), each with its own life cycle, but note that the overlapping life cycles generally match in their durations and events. The life-cycle colors indicate product availability and attendant support. Red indicates no availability, yellow indicates some availability, and green indicates high availability.

Each OT product life cycle starts with research and development, followed by a release to the market. While the product generation is available and supported, updates are produced by the vendor with new features or security patches. When a product generation is nearing its end of life, the vendor puts out a “last call” for product purchases, followed by an “end of support” announcement. Each product generation life cycle lasts about 10 years, with 4 to 5 years devoted

to R&D, and 5 to 6 years devoted to production and support. When one R&D cycle ends, the next R&D cycle begins to level the workload for product development and support and provide some predictability to customers. Note that gray market support of each product generation is shown in Figure 2-2, although such support should not be construed as part of the actual product life cycle.

Three different cyber security life cycles are shown at the bottom of Figure 2-2. The colors in each life cycle indicate a level of security. Green indicates acceptable security, yellow indicates weakening security, and red would indicate unacceptable security (not shown in Figure 2-2).

The first cyber security life cycle, indicated as “kill chain,” represents a simplified version of the various stages of a cyber attack. An attack begins with reconnaissance, followed by targeting and development (sometimes called *weaponization*), followed by delivery so the attacker can inspect the OT system, install the weapon, and erase any evidence of intrusion. A successful attack would ultimately launch the weapon and cause harm by the theft of data or manipulation of plant systems, resulting in facility or system damage, lost generation, or reputational damage. In Figure 2-2, the kill chains are terminated via detection and recovery methods and procedures used by the facility owner. Four different kill chain life cycles are shown in Figure 2-2. Note that they have the same durations when, in reality, kill chain life-cycle durations can vary from days, weeks, months, and years.

The second cyber security life cycle is indicated as “Control Method Persistence,” marked by emerging vulnerabilities that are mitigated via OT updates and refreshes. For this report, persistence means the ability of control methods to maintain security effectiveness over time. The security effectiveness of control methods may decay over time as technologies change and threat actors become more knowledgeable. Notice a vertical red line in Figure 2-2 at year N+1. The Control Method Persistence life cycle indicates some emerging vulnerabilities at year N (today), and the facility owner is aware of a Generation X product update under development by the vendor. In year N+1, the product update is staged and ready for installation via facility procedures, simply waiting for the first opportunity for an OT system outage. The owner is willing to operate the facility with some risk and is not willing to shut down the facility only for installing the update. However, for other reasons, the facility undergoes a curtailment order in year N+1, and the update is installed, restoring cyber security health to green in the control method persistence life cycle. Another update is shown in year N+11 in Figure 2-2, and in between updates are some emerging vulnerabilities mitigated via OT system refreshes in years N+4 and N+9. Note that control method persistence weakens in the last three years of the facility life because the owner is not investing in a refresh and is operating the obsolete Generation X+2 of the OT product with gray market support.

The third and last cyber security life cycle is indicated as “Threat Intelligence,” marked by external cyber events related to the OT technology in the owner’s facility. External event advisories indicate significant vulnerabilities, intrusions, attacks, data theft, and so on, with recommended corrective actions. When an external event occurs, relevant to the facility owner, a cyber security reassessment is performed to determine the extent and scope of new or modified security controls necessary to mitigate the attendant risks. Installation and/or adoption of new or modified controls follow each reassessment, restoring the threat intelligence life cycle to green. Note that cyber security events are not periodic and can occur at any time. For example, three events occur in succession in year N+14 in the threat intelligence life cycle, followed by one reassessment.

3

CONSEQUENCE ANALYSIS AND PRIORITIZATION: CRITICAL FUNCTIONS AND SERVICES THROUGH LIFE CYCLES

This section presents practical methods for discerning what is most important to the asset owner—and how those priorities change over the business and plant’s life cycle. Effective cyber security strategies no longer focus on how to avoid being “hacked”; instead, they acknowledge that the threat exists and must instead focus on identifying, triaging, and protecting what matters most to the owner.

While the critical functions and services of one fossil generation facility will not vary too far from those of another fossil generation facility, each owner’s threshold of acknowledged and accepted risk may fluctuate based on when in the life cycle the company, equipment, or facility is operating.

If the critical function and service of a fossil fuel power station is to produce electricity, a threshold must be identified to address “to what extent” an asset owner can go without producing electricity. Consequences other than loss of generation must also be considered, such as safety, equipment damage, reputation, and so on. Consequence analysis can describe a duration, impact, severity, or other quantifier that identifies a threshold that, if surpassed, would result in an unacceptable consequence.

An owner may analyze for different hazards with different consequences, such as the following:

- Lost generation
- Personnel safety
- Business reputation
- Chemical release

In this report, a hazard is a system condition, which if exploited, can lead to a consequence. A hazard analysis may identify different digital systems that contribute to different hazards and may yield different acceptable risk levels based on the hazard. It is important to note that both CCE and TAM assume that the analyzed hazard will occur if the system is successfully attacked, and therefore it requires mitigation of the attack surface to an acceptable risk level for analyzed hazards. In the case where a hazard analysis is not performed to determine the consequence, CCE and TAM recommend using the most conservative consequence level.

There are several hazard analysis methods for determining consequences that are described in this section. How CCE and the TAM use the output of CCE or other hazard analysis methods is described in Section 4. This section also explores how life cycles influence risk decisions using the first phase of CCE [3] to identify, examine, and prioritize HCEs.

A summary of CCE and TAM is provided in Section 4.

CCE Consequence Analysis Best Practices

In CCE, analysts call the defined critical functions and services the *objective*. The thresholds that define the systems and processes involved in the objective make up the *scope*. These definitions are paramount in CCE; however, when examining the objective and scope through a lens that adapts to the various stages of a life cycle, it is apparent that the owner-defined objective and scope must adapt through the life cycle to stay relevant to the owner's business needs.

If a fossil generation plant owner's critical function and service are to produce electricity for the region, the adversary's objective would be to stop, prevent, or degrade the plant's ability to produce electricity. The scope may set the threshold, which would cripple the plant to five days.¹ That is, if an *event* were to happen—cyber attack or otherwise—that prevented the plant from generating electricity for five days straight, the business would fail.

Note that it is much more likely that an adversary's objective will evolve through the plant's life cycles than it is that the specific events will. In other words, the adversary may need to only ruin a proof of concept or testing phase in early life, whereas plants close to decommissioning may require a much more substantial outage for the adversary to cause crippling damage to the plant. In both scenarios, an adversary may still attempt to disable or destroy turbines in the same way.

Also, consider how the plant's threshold for criticality varies, depending on what point in its life cycle it is. Preventing a generation plant from producing electricity in early, pre-commissioning phases would be inconsequential; the plant is not producing electricity anyway.

This helps illustrate why plant owners need to adjust their cyber security strategies and approaches temporally. Recognizing that a triage approach must adapt to the actual business needs—needs that also adapt continually—is imperative for effective cyber security.

The CCE methodology starts with brainstorming the potentially disruptive events that would significantly inhibit a generation plant's ability to provide the critical services and functions required to generate electricity. During previous CCE engagements in a variety of critical infrastructure sectors, a few key strategies have evolved from trends to assist asset owners in brainstorming possible attack methods (and their respective consequences). As a result, three targeting categories have emerged to help CCE practitioners explore adversarial approaches.

The first category of targets to consider is coined "physical infrastructure and interdependency." These are the *physical elements* used in a process function. In the electric sector, these may include generation, substation, or control center facility elements. When these physical elements are identified, look for interdependencies or choke points in the infrastructure.

¹ The five-day duration was assigned arbitrarily for the purposes of illustration. The actual threshold is defined by a governing body over the critical infrastructure and/or the process owner.

Example 3-1

Physical infrastructure and interdependency [5]

Infrastructure example:

Affecting cooling water supply system for a power generation facility with the intent to impact power delivery for all units at a station. Disruption of cooling water supply will force all units to shut down.

Interdependency example:

For an electric utility with assets that include gas-fired electrical power generation station(s), a “choke point” example would be the natural gas delivery system, most typically a pipeline infrastructure. The power generation plants are dependent on the natural gas delivery system and/or natural gas supplier (in the natural gas supply chain, this describes the natural gas producer, which can often be a company separate from the natural gas delivery/pipeline asset owner). The choke point could be targeted directly (delivery system or production system attack) or indirectly (attack on the asset owner of the delivery system or production system).

Methods of affecting infrastructure, as described in Example 3-1, could include targeting the overcurrent protection of physical assets and then misoperating devices to cause physical effects. Switchyards contain a wide variety of electrical infrastructure elements that can be misoperated to impact generation or other critical functions and services of a fossil power plant. These elements may include breakers, switches, transformers, protection relays, voltage load tap change, capacitor banks, and circuit reclosers [5].

The second targeting category is the horizontal application of technology. This examines targets that are widely deployed, either within the system or across a geographic region. Consider function-specific, widely deployed ICS technologies belonging to the same technology vendor platform, like vendor-specific implementation models of PLCs, remote terminal units, protection relays, meters, and so on. Often, single or even multiple instances/versions of these devices may be deployed throughout a critical infrastructure business enterprise for both geographically dispersed and localized asset models.

Another aspect to consider is the increased “depth” of a technology deployment; that is, there is an incentive to develop and adopt vendor solutions that integrate new and previously deployed, legacy technologies through common programming and monitoring applications. These legacy applications of technology may be especially common in plants nearing decommissioning. This broad and deep functional coverage within the systems is also attractive and valuable to a potential threat actor [5].

Example 3-2

Horizontal application example [5]

An electric utility may consolidate on a specific vendor's piece of equipment to drive consistency from site to site and reduce the level of system complexity for their field personnel. If a payload targeting the common device was deployed throughout the fleet through targeting and misuse of engineering or maintenance software/procedures, the corrective actions to repair/replace the compromised hardware would be extremely time-consuming, if not impossible, from a workforce perspective.

On April 3, 2020, a modified line of code in two fully redundant control centers simultaneously caused the loss of the same critical automatic generation control (AGC) task in both control centers. Because the same technology, firmware, software, and control systems were used to set generation schedules and automatically calculate area control error, both functionalities were lost to the utility [6].

Finally, the third category of targets looks at the increasing “reliance on automation and control capabilities.” Successful attacks against related targets would inhibit an organization’s automation or control functions. To achieve highly reliable delivery of services, there needs to be a system that can detect faults or system events and automatically respond or reconfigure to continue to provide services. Within most critical infrastructure organizations, there are systems and processes that have been automated in order to provide functionality that cannot be delivered manually with the necessary real-time response to ensure system reliability and safety.

Power generation facilities, regardless of fuel type, rely heavily on resource inputs like automated fuel management systems, feed water systems, water cooling systems, unit control systems, voltage regulation, and a wide variety of system protection controls that prevent damage or mitigate safety risks. Adversaries can target any one of these automated systems individually, or they may recognize the redundancies in place and choose to misuse or manipulate multiple systems simultaneously [5].

Example 3-3

Reliance on automation and control capabilities examples [5]

Consider the following examples in the “reliance on automation and control capabilities” category of targets:

- *Any “real-time” remote monitoring and/or control of assets*
- *Same day modifications to natural gas receipt and delivery volumes*
- *Timely collection of accurate volume, gas constituent, and operational parameter data in a geographically dispersed set of system assets*
- *Electric utility energy management system (EMS) and energy load balancing systems*
- *Power system area balancing through AGC and scheduling*

- *Power element maintenance ticketing and electronic-tagging systems*
- *Use of automatic load shedding schemes within the EMS (Special Protection Schemes, Remedial Action Schemes)*

Each of these target categories will help asset owners brainstorm and identify targets that an adversary can use to cause devastating consequences to the plant. These categories can also serve as talking points to guide necessary conversations at different stages of the plant's life cycle.

Does the plant require, use, or rely on the same automation capabilities now as it did before commissioning? Do operators and engineers rely on a control function that did not exist before implementing the technology? Have any changes in the plant's operating procedures during its current stage of the life cycle required (or eliminated!) a monitoring function?

The CCE methodology encourages and guides asset owners to ask the right questions during a consequence analysis. The emphasis on defining the critical functions early in the process assists scoping and limiting the cost of a robust but concentrated cyber security program. The three categories of targeting strategies help the CCE Team identify real targets that an adversary may use to destroy the organization's critical functions. Using CCE as a powerful tool for consequence analysis and prioritization, stakeholders can concentrate their efforts on securing what matters most.

TAM Consequence Analysis Best Practices

The TAM uses the output of CCE or other consequence analysis in a risk informed process to determine the level of mitigation required for an attack surface. A consequence analysis output is used by the owner to set a TAM consequence *Target Level Threshold* for a *Combined Security Effectiveness* score for each exploit sequence. When the Combined Security Effectiveness score for allocated control methods for each exploit sequence reaches the target level threshold, the exploit sequence is mitigated to an acceptable risk level, as determined by the owner. When all exploit sequences are mitigated to the target level threshold, the device or system attack surface is fully mitigated.

A summary of the TAM's risk informed approach to attack surface mitigation is provided in the *TAM Risk Informed Attack Surface Mitigation Best Practices* section.

EPRI has published a separate *Risk Informed Target Level Topical Guide* (3002015760) [7] that provides more detailed guidance on how to use consequence analysis and target level thresholds in the TAM.

Other Consequence Analysis Methods

In addition to CCE, there are formal and informal consequence analysis methods available to facility owners. Three hazard analysis methods are described in the following:

- *EPRI Guideline for Establishing Criticality Rankings for Fossil Power Plant Systems and Components* (1025341) [8]
- IEC 61508 Functional Safety of Electronic Systems and 61511 Safety Instrumented Systems (SIS) [9, 10]
- IEC 62443 Security for Industrial Automation and Control Systems [11]

EPRI Guideline for Establishing Criticality Rankings for Fossil Power Plant Systems and Components (1025341)

EPRI Guideline for Establishing Criticality Rankings for Fossil Power Plant Systems and Components (1025341) [8] was developed as part of the Equipment Reliability (ER) program in the EPRI Generation sector. The guideline describes three component criticality ranking (CCR) approaches from low complexity to high complexity, as well as a method to assign system criticality rankings (SCRs).

Low Complexity CCR Method

The low complexity CCR method uses experienced plant system engineers and/or operations personnel to review the master equipment list and make an ER classification for each piece of equipment based on the system engineers' and/or operators' experience and knowledge of plant systems and equipment functions. There are three ER component classifications. *Critical component* represents a critical component to maintaining an important system function. *Non-critical* represents a non-critical component that is not critical to sustaining generation, safety, or environmental compliance, but could result in a costly repair if failure occurs. *RTF* represents a run-to-failure component that is not critical to sustaining component or system functions, nor would failure result in costly repairs.

Medium Complexity CCR

The medium complexity CCR method establishes equipment criticality, along with an operating classification through a series of 14 equipment criticality determination questions and is often performed as part of a maintenance basis optimization process. The results of the method are nine Importance Values with Importance 1 through 5 for critical equipment and Importance 6 through 9 for non-critical equipment (see Table 4-1 in the CCR report).

High Complexity CCR

The EPRI CCR report describes high complexity approaches that apply multiple weighting factors, input variables, and algorithms to produce more quantitative and granular assessments of CCR. The report uses the System and Equipment Reliability Prioritization (SERP) process as an example. The SERP process also uses importance measures based on a common set of factors across generating facilities (for example, safety, environmental, cost, and reliability) that result in scored ranking. The rankings are established as consensus input from operations, maintenance, technology specialists, craft, and management.

The SERP process ranks each individual system to generate an SCR value for each system. The SERP process also ranks the importance of each piece of equipment to the function of the system, which results in an operational criticality ranking (OCR) value. The overall CCR value is established as a product of SCR and OCR, as follows:

$$\text{Component Criticality (CCR)} = \text{System Criticality (SCR)} \times \text{Operational Criticality (OCR)}$$

In the SERP example in the report, the CCR values range from 88.22 to 274.81.

System Criticality Ranking

The EPRI CCR report also provides a method to assign criticality rankings to systems. The CCR system ranking is composed of four tiers, Tiers 1–4, where Tier 1 is critical to generation or some other hazard. The process used to establish the tiers is similar to the Low Complexity CCR method by bringing together plant engineering, O&M leadership, and other plant knowledgeable individuals to assign a value to each system.

IEC 62443 Security for Industrial Automation and Control Systems

The process industry is developing a series of cyber security standards published in four parts and multiple subparts in IEC 62443 Security for Industrial Automation and Control Systems [11]. The four parts of IEC 62443-1 provide general concepts, concepts and models, and a glossary of terms and conditions. The four parts in IEC 62443-2 provide policies and procedures, whereas IEC 62443-3 has three parts associated with system level topics, including Security Levels (SL). The two parts of IEC 62443-4 are related to supply chain security for system life cycle and component requirements.

Similar to Safety Integrity Levels (SIL) described later, IEC 62443-1-1 defines SLs that determine the level of required security, and up to seven functional requirements to be applied to a system based on the means and resources needed to compromise the system. IEC 62443-3-2, Security Risk Assessment and System Design (in draft) applies Target Security Levels (Target SL) to zones and conduits within the system under assessment. A risk assessment is performed to determine the appropriate Target SL. A specific risk assessment method has yet to be developed; however, if a facility owner has performed a risk assessment and assigned Target SLs, then Target SLs can be mapped to TAM Target Levels.

Current practice includes five Security Levels, SL 0–SL 4, with SL 4 as the highest security level.

IEC 61508/61511: Functional Safety of Electronic Systems

The process industry has developed a process hazard analysis approach that is described in IEC 61511 Functional Safety – Safety Instrumented Systems for the Process Industry Sector [10], published in three parts, and IEC 61508: Functional Safety of Electrical/Electronic/Programmable Electronic Safety-Related Systems, published in seven parts [9].

Hazards in process plants are mitigated by installing safety instrumented systems (SIS) as part of a multi-layered safety approach. SIS are installed in addition to the normal I&C systems where the hazard analysis indicates that additional risk reduction is required. The SIS take the process to a “safe state” when calculated safe operating conditions have been exceeded.

A process hazard analysis (PHA) and layer of protection analysis (LOPA) method is used to analyze all known process hazards and layers of protection. The LOPA analysis determines

whether existing control system reliability and independent barriers are sufficient so that risk is acceptable. If not, then an SIS is required. Each safety function to be implemented via a SIS is assigned a target safety integrity level (SIL).

SIS reduce risk by implementing safety instrumented functions (SIF) to achieve or maintain a safe state. A Target SIL level for each SIF is determined by using a risk reduction factor to reduce the hazard frequency to an acceptable level. There are four SIL levels, 1–4, with 4 representing the highest risk reduction. SIL 3 is the highest safety integrity level that is economically feasible for most process facilities. Therefore, the practical result of applying the IEC 61508 and 61511 hazard analysis is installing a SIS with safety functions of SIL 1, 2, or 3 SIS, or no SIS is required.

Consequence-Based Critical Function Prioritization Using CCE

This section articulates how CCE Phase 1 (Consequence Prioritization) can be applied in the fossil power industry. It also examines how the identification and prioritization of consequences evolve through the various stages of the refined life cycles discussed in the previous section. It presents case studies to discern how consequence prioritization changes the way generation facility owners identify and prioritize critical functions, given limited resources.

The goal of CCE is to protect critical functions—those services that define an organization’s purpose—from cyber threats while proactively preparing for the next generation of cyber attacks. CCE is a triage approach that identifies an organization’s most critical operations, addresses what really bad days (worst-case scenarios) or HCEs look like for that specific organization, determines which of those worst-case scenarios would be the most crippling, and seeks to protect the organization from those specific HCEs through mitigation.

The CCE construct consists of four phases. An organization moves linearly through each phase to brainstorm events, protect data, perform a critical analysis, identify risks, and establish mitigations. The four stages of CCE include the following:

- Phase 1: Consequence Prioritization
- Phase 2: System-of-Systems Analysis
- Phase 3: Consequence-based Targeting
- Phase 4: Mitigations and Protection

Just as performance, profit, repairs, outages, and market prices vary greatly across the different stages of facility and business life cycles, so do the critical functions and threats of those devastating, irreparable HCEs. HCEs refer to those adverse cyber events deemed most significant in terms of the severity of impact. Organizations define criteria to help rank HCEs in a manner that best represents each organization’s mission and business values. The argument that these criteria must evolve through the facility’s life cycle is made in this section.

The flexibility in HCE prioritization illustrates how CCE can adapt to different stages within a facility life cycle. For example, the cost of recovery and the severity of damage caused by cyber-enabled sabotage may be the paramount criteria when weighting HCEs early in the facility life cycle, but those same costs may not predominate concerns during end-of-life operations. Let's revisit the case study discussed in Section 2 to illustrate how risk tolerance at end-of-life can affect consequence prioritization.

Example 3-4 **Facility near the end of life**

Since the facility will be decommissioned within the next four years, the owner begins lowering his risk tolerance. As the owner limits funding for maintenance and repairs for most systems, he reassigns his priority to the systems that preserve operator safety. Further, the owner understands the risk—if the systems fail before end-of-life and are too costly to repair, he will terminate the entire plant's operations.

This shift in risk tolerance demonstrates the urgency to reduce operating costs through the end of life. This shift also illustrates how an attack that cripples operations—while still devastating and would likely result in early decommissioning of the plant—may not be as big of a concern to the owner during this stage of the facility's life cycle. This type of attack would have a much more irreparable effect at an earlier stage in the plant's life cycle.

This point is strengthened by the owner's pledge to delay or avoid the installation of security patches or cyber security upgrades. The owner accepts this risk and would rather allocate his resources and time to other areas of operations.

Introduction to Hypothetical HCEs

The primary goal of Phase 1 of CCE is to identify potential disruptive events that would significantly inhibit an organization's ability to provide the critical services and functions deemed critical to their business mission. These events can be realized through scenarios that involve attacks on critical infrastructure owner assets.

Below are two examples of “really bad days” that can both hypothetically happen at a generation plant and are possible to achieve through cyber means. These scenarios were developed by senior analysts with significant input from engineering and sector-specific SMEs as part of a larger CCE project.

In a typical CCE engagement, these HCEs would be scored and weighted against criteria defined by the asset owner to prioritize mitigation efforts in later phases. However, for the purposes of this effort, prioritization emphasis will be placed on the different stages of the plant's life cycle.

Just as the threshold for the criticality of consequences changes through the plant's life cycle, so do the TAM exploit sequences identified in TAM Step 1. Consequences to the owner's critical functions vary with plant age; similarly, asset features and data flow manifest differently in the various stages of the plant's life cycle. This illustrates how TAM Step 1 should be executed concurrently with CCE Phase 1 at each stage of the fossil generation facility's life cycle. Section 4 continues the discussion on best practices for managing life cycle changes and risks.

HCE #1—Lubrication Oil System Failure

This example considers how cyber-enabled sabotage against a fossil generation facility's lubrication oil system can have catastrophic effects in the form of extensive damage to a steam turbine and lost production. There are several ways that cyber-enabled sabotage against the lubrication oil system could affect a fossil generation facility's critical function of generating power (inadequate cooling, overloading, overheating, and so on). Still, this example focuses on an attack that causes a loss of oil.

A properly working lubrication oil system ensures the health of an engine and reduces the stress of normal wear and tear. Like any engine, turbine bearings in fossil generation facilities require lubrication. In a turbine, the lubrication oil system supplies oil to the journal bearings and creates an oil wedge that reduces friction and helps maintain lower system temperatures [12]. In hydrogen-cooled generators, the oil also serves as a sealing medium to prevent the loss or leak of hydrogen along the shaft [6].

The lube oil system includes pumps, motors, controls, piping, valves, flow control devices, filters, and coolers. Protective schemes are also incorporated to protect the equipment against low lube oil supply, low pressure, and off-specification temperatures [13]. A failure of the lubrication oil pump may cause the turbine bearings to seize. Seized turbine bearings may cause significant and costly damages to additional components.

This hypothetical is not too improbable. The International Association of Engineering Insurers published that a loss of oil is the number one culprit behind steam turbine failures worldwide [12]. An adversary may certainly see the incentive to sabotage the lubrication oil system. A successful attack not only forces repair damages, but it may also cost the turbine owner significant revenue in terms of lost generation time.

Example 3-5 Lubrication oil pump failure

At the commissioning of an unspecified plant, a power interruption resulted in a loss of the control system. The plant had three combustion turbines (375 MW) and planned the construction of a 178-MW steam turbine to allow the plant to operate in combined-cycle mode. As a result of the loss of power and the resulting loss of the DCS, the auxiliary oil pump did not start after the trip. An emergency pump also did not start after the trip, and all lube oil was lost during roll down.

The damage to the steam turbine was extensive. The attack included damage to the bearings, rotor, inter-stage seals, and the blade. It took months to replace the turbine and support systems on a skid, especially since the mechanical damage extended beyond the turbine housing.

The total financial loss was \$12 million in repairs and \$30 million in lost income.²

² This example was created by Senior Analysts Sarah Freeman and Colleen Glenn of Idaho National Laboratory on October 5, 2016. They cited the following work: Wallace Ebner, "Strategies for the Prevention of Turbine Lube Oil System Failures," in *Proceedings of the ASME 2013 Power Conference*, July 29–August 1, 2013, Boston, MA.

This example cost the plant owners lost revenue, significant damages, considerable repair, and an extensive amount of time. This is certainly a worst-case scenario; however, the plant's age may either mitigate or compound the extent to which this HCE affects its critical functions. If we break a plant system's life cycle into only three stages for the purposes of discussion, we can observe how this specific HCE incapacitates operations to different degrees. Figure 3-1 illustrates how the consequences of an attack that disables the auxiliary oil pump change, depending on the facility's life cycle.




Figure 3-1
Image comparing the impact of HCE #1 on a facility during three different stages of its life cycle

Lost revenue is almost always a significant concern for any business. Although the HCE may delay commissioning so long that it affects revenue, the amount of lost revenue is greater during mid-life operations (year N to N+10 of the Plant System Life Cycle in Figure 2-2). If the attack occurred during end-of-life operations (year N+10 to N+15 and beyond of the Plant System Life Cycle in Figure 2-2), the value of the lost revenue is the same as mid-life operations; however, the lost revenue during this phase is often unreliable income. That is, a loss of this income is anticipated to be an assumed risk of acceptable loss, and such loss would not be catastrophic for the plant.

As mentioned, if the HCE occurred during end-of-life operations, it will almost certainly result in a decision to retire or decommission the facility early. On the other hand, if the HCE occurred during pre-commissioning, it *may* delay operations. An attack on mid-life operations delays production until repairs are completed.

Last, the costs associated with repairing the damage to the bearings, rotor, inter-stage seals, and the blade remain the same. However, the \$12 million in repairs is required in early-life and mid-life operations. That cost can be ignored in end-of-life operations if the asset owner decides to decommission the plant.

The HCE in Example 3-5 illustrates how the consequences and impact of an attack change throughout a plant's life span. Regardless of the life-cycle stage, an attacker gains access to a network, disables the power, overrides the auxiliary oil pump, and causes significant, disabling damage to the turbine. The *consequences* of the attack, however, vary greatly depending on the plant's age. This helps illustrate how cyber risk must be managed differently within the context of time. In other words, the facility's age should help steer its cyber security strategy, just as it does its risk tolerance.

HCE #2—Communications System Disruption

Consequences in CCE and TAM do not always involve fatalities or physical damage to the equipment or facility. Sometimes, harming a plant's reliability or efficiency can be catastrophic to its reputation. This loss of system confidence can, in some cases, be more debilitating than the costs associated with the repair.

An attack designed to undermine the public's, government's, or investors' confidence in the plant or business is a real threat. The prevention of such attacks should be a driving effort when developing strategies to mitigate vulnerabilities during each plant's life-cycle stage.

The supervisory control and data acquisition (SCADA) system, I&C equipment, and control automation systems are just a few examples of systems at fossil generation facilities that rely heavily on network communication. Plant owners are increasingly implementing equipment to control reliability and efficiency, as well as provide remote monitoring capabilities to process owners. All these capabilities both require and rely on constant access to real-time data [14].

If the communication infrastructure used by these automation, monitoring, and remote-control systems is disabled, operators and owners will lose the ability to monitor and control critical functions and capabilities. The plant would be forced to operate in manual mode or halt all operations until restoration.

Example 3-6 **Malicious load shedding in Ukraine**

In December 2015, an attacker targeted the distribution management system of three utilities, opening breakers and shedding load, which resulted in an outage for 225,000 endpoints [13]. Restoration costs following the event for one utility have already totaled \$300K, and confidence in system integrity remains an issue.³

Fortunately, for the Ukrainian utility owners, the outage was quickly resolved because their electric grid was almost completely recoverable in manual mode. What is interesting in this case, and certainly applicable to the fossil generation facility owner, is the damage to system integrity confidence.

As previously mentioned, any automation functionality in fossil generation facilities relies on communication. If this communication is disrupted, the functionality is lost. One example of a cyber attack that targets the communication infrastructure is bricking.⁴ A bricked device is one that is permanently unfunctional and often not recoverable. To worsen the situation, a cyber attack involving bricking is often accompanied with widespread malware attacks to other devices. Consider a theoretical example of bricking that would cripple the confidence in a fossil generation facility's system integrity.

³ This example was captured by Senior Analysts Sarah Freeman and Colleen Glenn of Idaho National Laboratory on October 5, 2016. They cited the work in this footnote to SANS.

⁴ Bricking is an emerging cyber attack increasingly used in the hardware hacking community. The attacker's goal is to render an electronic device useless.

Example 3-7

Rendering communication systems inoperable

After gaining access to networked communication devices, an adversary pushes a malware-laden firmware update that bricks multiple network assets, rendering them permanently non-functional. Additional networked devices are infected with malware. Digital investigation and forensics discover the extent and boundaries of affected network devices, but this takes a considerable amount of time to determine. Recovery is complicated; operators need to configure network devices from scratch.

Recall in Example 3-5 (lubrication oil pump failure) that the repair and recovery costs can be compounded or even ignored, depending on the facility's age. Example 3-7 would likely include the monetary cost of replacing bricked devices; however, for the purposes of examining how consequences to a plant's critical functions change over the life cycle, we will focus on the loss of confidence in system integrity in this section.

System integrity confidence describes whether restoration and recovery efforts can restore system integrity with confidence following an adverse event. After the network engineers configure network devices from scratch and operators replace bricked devices, will the system operate as expected? Did the initial attack propagate in multiple systems and confound restoration efforts? (In this example, it did.) Will malicious operation by unauthorized users persist after recovery? These concerns illustrate intangible ways an organization's confidence in their system may be negatively impacted after a cyber attack.

Figure 3-2 illustrates how the consequences of a cyber attack with regard to the system integrity confidence criterion should be weighted more heavily in early stages of plant life than mid-life or at decommissioning.




Figure 3-2

Image comparing the impact of HCE #2 on a facility during three different stages of its life cycle

Before commissioning a plant, operators and engineers will not have the same confidence in the process as they would in year N to N+10 of the Plant System Life Cycle in Figure 2-2. Although a complete loss of system integrity confidence is unlikely at end-of-life stages of the facility's life cycle, such loss of confidence is a completely irrelevant consequence of a cyber attack right before facility retirement.

4

BEST PRACTICE STRATEGIES FOR CYBER ENGINEERING

The previous sections discussed the framework that shows the temporal nature of generation plant life cycles. They identified that refined life cycles for the business, vendors' businesses, facility, plant systems, components, and digital products are broken into several stages, each stage with its own specific performance and risk issues. As such, cyber risk can only be properly and efficiently addressed with decision-making that places emphasis on the balance between the lifetime and the specific point in time.

This section identifies and recommends best practice strategies in the context of an overall life cycle framework. The strategies in this section were developed based on the plant life-cycle characteristics described in Section 2, as well as the consequence analysis and prioritization explored in Section 3.

This section aims to provide best practice *strategies* and best practice *approaches* to mitigating cyber security risk for fossil energy facility owners. The section does *not* provide any specific software platforms or tools to mitigate risk. As demonstrated in the previous sections, consequences, priorities, critical functions, and cyber security threats evolve through the different stages of the plant's life cycle. No single cyber security tool can properly protect a fossil generation facility's most critical functions from pre-commissioning to decommissioning.

Protecting ICS, or OT, is challenged by the lack of a deliberate approach, preventing a complete understanding of the systems' characteristics. Adversaries are deliberate in understanding system characteristics to a deep technical level, enabling effective attacks and compromises. Consequently, defenders must understand their critical systems as well as or better than adversaries to provide effective and efficient protection.

Successful best practice strategies and approaches are consistent with the first principles of OT cyber security. First principles have been described in philosophy, logic, and physics: "In physics, a calculation is said to be from first principles, or *ab initio*, if it starts directly at the level of established laws of physics and does not make assumptions such as empirical model and fitting parameters" [15]. Similarly, cyber security and cyber security engineering and design must start from a fundamental set of truths about cyber security for systems engineering.

A cyber security program should consist of risk-informed systems engineering approaches that are built upon and adhere to first principles—like CCE and TAM—working in concert to bound the attack surface and adapt to the changing threat landscape through each phase of the life cycle. These systems engineering risk-informed *approaches* provide the plant owner with the affordable complaint, thorough, and versatile best practices required to adapt to the changing threat landscape through each stage of the plant's life cycle.

Various standards, regulations, and guidance are available to assist a plant owner in developing an effective cyber security program. However, the industry is recognizing the need to begin from first principles using a system engineering and risk informed approach. CCE and TAM are consistent with other guidance and standards because CCE and TAM are built upon the first principles. Some examples of other cyber security standards and guidance are the following:

- NIST SP 800-53 Rev. 4: Security and Privacy for Federal Information Systems and Organizations
- NIST SP 800-82 Rev. 2: Guide to Industrial Control Systems (ICS) Security
- NIST Framework for Improving Critical Infrastructure Cyber Security, V 1.1
- NIST SP 800-60 Vol. 2: Developing Cyber Resilient Systems: A System Security Engineering Approach
- IEC 62443, a series of standards for control cyber security management

This section is organized as follows:

- Cyber security first principles
- Summary of CCE and TAM as first principle best practices
- Best practice strategies, and how CCE and TAM and other selected guidance implement the best practice strategies, as follows:
 - Using the output from hazard and consequence analysis
 - Identifying attack surface vulnerabilities
 - Risk informed attack surface mitigation
- CCE and TAM case study examples

Cyber Security First Principles

A first principle is a foundational proposition or assumption that stands alone. We cannot deduce first principles from any other proposition or assumption.

Aristotle, writing on first principles, said the following:

“In every systematic inquiry (methodos) where there are first principles, or causes, or elements, knowledge and science result from acquiring knowledge of these; for we think we know something just in case we acquire knowledge of the primary causes, the primary first principles, all the way to the elements” [16].

Later, he described how knowledge is derived from first principles, defining first principles as “the first basis from which a thing is known” [17].

Elon Musk has famously used first principles in developing his space travel strategy and SpaceX:

“I tend to approach things from a physics framework,” Musk said in an interview. “Physics teaches you to reason from first principles rather than by analogy. So, I said, okay, let’s look at the first principles. What is a rocket made of? Aerospace-grade aluminum alloys, plus some

titanium, copper, and carbon fiber. Then I asked, what is the value of those materials on the commodity market? It turned out that the materials cost of a rocket was around two percent of the typical price” [18].

Instead of purchasing space vehicles for tens of millions, Musk decided to start his own company, purchase the raw materials at a fraction of the price, and build the rockets himself. SpaceX was the result.

Within a few years, SpaceX had cut the cost of producing and launching a rocket by nearly 10 times while still making a profit. Musk used first principles thinking to break the situation down to the fundamentals, bypass the high prices of the aerospace industry, and create a more effective solution [5].

First principles thinking is the act of boiling a process down to the fundamental parts that you know are true and building up from there.

OT cyber security best practices are built upon fundamental OT systems engineering first principles, risk management first principles, and a set of engineering and design first principles. The first principles are listed below. All effective best practices incorporate these first principles and are regulatory agnostic and future proof.

Systems Engineering First Principles of OT Cyber Security

Systems engineering is the foundation of OT system design, operation, and maintenance. Understanding how OT cyber security is tied to systems engineering is imperative for any best practice to be effective.

Principle 1 (SE1): OT cyber attacks can have physical consequences.

Industrial control systems control physical processes. The process comprises components that operate together as a system with a defined purpose. The systems function together to serve a business purpose.

Principle 2 (SE2): The adversary seeks to disrupt, disable, destroy, or control digital systems that result in a negative consequence.

A cyber attack targets an owner’s digital components and/or systems for the purpose of disrupting, disabling, destroying, or maliciously controlling a digital system (or destroying the integrity of the data or stealing critical data), resulting in a negative consequence (impact).

Principle 3 (SE3): Successful attacks require an existing pathway and exploit mechanism.

A successful attack must use a known attack pathway and exploit mechanism (for example, firmware supply chain to introduce altered firmware, controller faceplate to alter controller configuration, network connection to introduce a denial of service [DOS] attack) in order to achieve an exploit objective. There is a finite number of existing attack pathways and exploit mechanisms that comprise the attack surface.

Principle 4 (SE4): Organizations must understand their own systems.

The business must fully understand their systems in order to effectively characterize and mitigate the attack surface.

Risk Management First Principles for OT Cyber Security

Cyber security risk is always present to some degree. The business must understand and mitigate cyber security risk to an acceptable level.

Principle 1 (RM1): Cyber risk combines an attack's probability (likelihood) and consequence (impact).

All risk is measured as a probability and consequence. Risk decisions are based on the ability and cost to lower the likelihood for a given consequence to an acceptable level, or the ability and cost to change the design to lower the consequence for a given likelihood to an acceptable level.

Principle 2 (RM2): Cyber risk changes over time.

Adversary capabilities change the likelihood of an attack over time. Aging technologies and deployment of new technologies (including software updates) may change the likelihood or impact over time.

Principle 3 (RM3): Cyber risk cannot be eliminated.

Cyber risk can approach zero but cannot be zero if digital devices are used to support business goals.

Principle 4 (RM4): Cyber risk can be mitigated.

Both the impact and likelihood can be reduced by applying design changes or control methods that fit within the business model and reduce risk to a level that is acceptable to the business.

Cyber Security Engineering and Design First Principles

Cyber security for OT systems is not a separate engineering activity. Cyber security must be incorporated in OT system engineering and design. Cyber Security Engineering and Design First Principles are derived from the System Engineering and Risk Management First Principles that allow for practical development of cyber security best practices.

Principle 1 (CE1): Security and safety are the same thing: they share the same hazards but with different causes.

Principle 2 (CE2): When threat capability is appropriately addressed, then an adversary's desire and motivation become inconsequential.

Principle 3 (CE3): Vulnerability assessment and design are driven by system design optimization. The granularity of analysis should be limited while maintaining boundaries between functional elements.

Principle 4 (CE4): A vulnerability consists of the attack pathway, the goal of the adversary, and the mechanism to achieve the goal through the pathway.

Principle 5 (CE5): Successful vulnerability assessments and design (mitigations) are bounded when:

- The attack surface has been reduced to its smallest scope.
- The remaining vulnerabilities have been identified.
- Design and implementation requirements, being informed from security control methods, have been developed and approved.

Principle 6 (CE6): At each level of system design, cyber security is effectively incorporated when:

- Vulnerabilities have been identified.
- Security control methods that protect, detect, respond, and recover from those vulnerabilities have been implemented.

Principle 7 (CE7): The three cyber risk elements (likelihood, vulnerability, and hazard) determine the risk acceptance level that informs the security control methods used to protect, detect, respond, and recover.

Principle 7a (CE7a): The likelihood for a successful attack (1) is based on the characteristics of both the vulnerabilities and the attacker's capabilities and (2) increases as the numbers of vulnerabilities and the attacker's capability increase.

Principle 8 (CE8): Cyber security is an integral part of engineering and system design. It starts once system functionality is first determined and continues until the design is complete and implemented.

CCE and TAM as First Principle Best Practices

Both the INL CCE and EPRI TAM approaches were developed from the cyber security first principles.

CCE Overview

CCE is a strategic methodology that prepares critical infrastructure organizations for existing and emerging cyber threats. Developed by INL, CCE provides both public and private organizations with a systematic approach to protect their most critical functions with cyber-informed risk mitigation strategies. This methodology provides organizations with an acute understanding of their system, a triaged list of solutions for the organization's worst-case scenarios, and an opportunity to demonstrate their commitment to securing the nation's critical infrastructure.

After extensive experience performing cyber security assessments on ICS, INL experts realized there needed to be a fundamental shift in how the nation handles control systems cyber security. INL leverages a unique combination of cyber researchers, control systems cyber security analysts, and control systems SMEs to assist national critical infrastructure organizations in protecting their critical operations from well-funded, highly skilled, and determined adversaries.

CCE is an ongoing engineering effort to eliminate inherent “trust assumptions” and mitigate cyber security gaps through calculated changes in technology, processes, and procedures, all aimed at “engineering out” vulnerabilities created by misplaced or unverified trust (Figure 4-1). CCE starts by reviewing the organization’s entire operation to identify the most critical systems and functions. Then, the organization performs a focused assessment on these most critical systems and functions. CCE does not just look at a vulnerability, but also it examines how an exploited vulnerability can result in costly or dangerous, physical consequences.




Figure 4-1
High-level overview of the four phases in the CCE methodology

The CCE methodology is iterative. It consists of the following four phases designed to guide the organization from identifying critical functions and services to engineering out the vulnerabilities that leave them exposed.

- Phase 1, *Consequence Prioritization*, identifies the organization’s most critical functions and services. Worst-case scenarios are brainstormed and prioritized based on their respective impacts to the organization, and HCEs are identified.
- Phase 2, *System-of-Systems Analysis*, identifies and collects organizational information on the critical systems involved in each HCE.
- Phase 3, *Consequence-based Targeting*, develops the paths, targets, access, and information—from an adversary’s perspective—required to cause the HCE to come to fruition.
- Phase 4, *Mitigations and Protections*, focuses on implementing mitigation efforts to prevent or detect the adversary’s ability to carry out the HCEs.

EPRI TAM Overview

EPRI has developed a risk-informed cyber security engineering approach that provides an objective technical approach for determining device or system attack surface and mitigations: *Cyber Security Technical Assessment Methodology, Risk Informed Exploit Sequence Identification and Mitigation, Revision 1* (3002012752) [19]. The need for the TAM was identified by utility members who required a more systematic and risk-informed approach that is cost-effective, cyber-secure, and regulatory-compliant. The TAM provides a different approach than the traditional security control requirements checklists, which are often regulatory-based and may require an evaluation of all controls regardless of efficacy or actual attack surface characteristics.

During the research to develop the TAM, it became clear that any cyber security approach must be developed from the cyber security first principles. The TAM is based on the first principles and guides the user through a methodical process that efficiently converges attack surface analysis and mitigation activities to an effective risk-informed result.

The TAM is an engineering process that examines the technical composition of devices and systems more closely and only applies those security control methods that are most effective, while considering potential plant consequences and exploit difficulty—the risk component—for the actual attack surface of the asset.

The process is efficient and complete by performing the following:

- Developing a baseline system/device configuration that provides a bounded starting point and reduces the attack surface as low as possible while still achieving functional objectives
- Narrowing the attack surface to the exploit objectives and attack pathways that are present
- Identifying the exploit mechanisms that remain to form exploit sequences that are then used to select appropriate mitigating control methods
- Using a risk-informed approach that incorporates consequence from hazard analysis and likelihood of a successful attack. These methods include intrinsic device, system, and program features and capabilities that can be leveraged to mitigate an exploit sequence and raise the exploit difficulty to a sufficiently high level
- Applying the best control methods based on the control methods' effectiveness and implementation burden
- Scaling to a full plant by grouping digital components and systems into relationship sets, where relationships exist, that are used to identify and allocate inherited shared control methods over entire systems where appropriate

The TAM is a three-step methodology, as shown in Figure 4-2.




Figure 4-2
TAM overview

Step One: Characterize the Attack Surface

Characterize the attack service by identifying all of the exploit sequences—based on the installed configuration and data flow—and documenting the results in a *Cyber Security Data Sheet* (CSDS). This approach is gaining traction across the industry as a standard way to communicate the attack surface.

Step Two: Identify, Score, and Allocate Engineered Control Methods

In Step 2, engineered control methods (those that are installed on the device) are identified, scored for security effectiveness, and allocated to exploit sequences—based on efficacy—until a consequence target level threshold value is achieved.

Any hazard can be analyzed—such as lost generation, radiological release, or reputational damage—to set targets for the level of consequence associated with that hazard.

The control methods installed on a device may not fully mitigate the entire attack surface (all exploit sequences). Exploit sequences that are not fully mitigated are referred to as *residual exploit sequences* and are addressed in Step 3.

Step Three: Mitigate Residual Exploit Sequences with Site- or Facility-Level Control Methods

Site control methods, like the physical security perimeter, are identified and scored in a shared control method library (CML). They are inherited through relationship sets documented in *Relationship Set Data Sheets* (RSDS) and allocated to mitigate the residual exploit sequences from Step 2.

Additionally, the steps taken to mitigate the attack surface are easily mapped to a regulatory compliance scheme, such as NERC-CIP, and can be documented in a regulatory requirements compliance map. The TAM identifies all the attack surfaces and mitigates them, and there is minimal, if any, impact from changing regulations because regulatory changes do not alter an attack surface.

Bounding the Universe of Software and Hardware Vulnerabilities

Regardless of any software or hardware vulnerability, an attacker still has to use one of the exploit sequences to achieve an exploit objective. Rather than attempting to mitigate all of the known and unknown software and hardware vulnerabilities (flaws and weaknesses) in a device or system, the TAM mitigates the exploit sequences. Thus, mitigating all exploit sequences bounds the universe of vulnerabilities.

The TAM also provides an approach to consider threat capability intelligence, that might include new exploit sequences or control method security effectiveness degradation.

Mitigation of the Entire Kill Chain

An interesting and powerful outcome of implementing the TAM across a system or facility is mitigation of the entire kill chain. Through aggregation, each CSDS inherits all of the exploit sequences and control methods (both engineered and shared) of all other CSDSs in a relationship set. This is an implicit inheritance that can be visualized in a data flow diagram. This ensures that all relevant kill chains are mitigated across all attack pathways without excessive expansion and redundancy in the CSDS.

TAM Implementation Use Cases

The TAM has been designed to be flexible and used in multiple use cases. An owner has the flexibility to implement the TAM as needed and does not need to implement the TAM for the entire plant from the start. All three use cases take advantage of the TAM's modularity and efficiency.

- TAM Use Case 1: Integrated into the overall digital engineering modification process.
 - As new digital systems, assets, and services are being considered and designed to be used in a critical infrastructure facility, the TAM provides detailed analytical information needed to assist an engineer with making well-informed decisions for mitigating cyber hazards.
- TAM Use Case 2: Use throughout the supply chain.
 - It is important to note that a vendor can prepare a “Reference CSDS” to deliver to a customer. The modularity of the TAM and its documentation artifacts allow it to be easily integrated throughout the supply chain, clarifying the division of responsibilities between the buyer and supplier, and reducing cyber hazards.
- TAM Use Case 3: In situ assessments as needed for components, systems, and services already installed.
 - The TAM can be implemented as needed for existing components and systems.

The Complementary Relationship of CCE and TAM

EPRI's TAM is a methodical process that efficiently converges the assessment and mitigation activities in a cyber security strategy to an effective result. The TAM baselines identify, assess, and mitigate cyber security vulnerabilities and exploits in a three-step, risk-informed approach that also reduces the burden of additional assessments.

While many consistencies exist between TAM and CCE methodologies, the implementation of both in concert yields the most value at every stage of the plant life cycle. For example, a significant product of the CCE process is the generation of the consequences stemming from cyber-enabled sabotage—the very same consequences that, when combined with hazard information, greatly complement the TAM process with hardly any extra added effort. The consequence and hazard information obtained through CCE can be combined with TAM results to make risk-informed cyber security decisions.

As mentioned previously, there is much value in implementing transformative approaches to a fossil energy plant's cyber security strategy. While positive and thorough cyber hygiene provides a solid foundation for the plant, it is alone an unsustainable strategy. Cyber hygiene is not enough to carry the process owners through the next generation of cyber-physical attacks.

Best Practices for Identifying Attack Surface Vulnerabilities

After a component or system has been prioritized for protection through consequence analysis, the attack surface needs to be characterized. If an attack surface is properly characterized, then allocation of security control methods to mitigate the attack surface can be straightforward. For an efficient and repeatable process, a common, bounded definition for vulnerabilities that allows for a consistent and straightforward approach for an attack surface analysis is needed. There are multiple definitions of an attack surface that usually include identifying vulnerabilities based on various approaches that range from purely academic theoretical models to specific software code weaknesses. However, attempting to characterize an attack surface through vulnerabilities and threats is an unbounded condition.

The term *vulnerability* in cyber security has several meanings. Past EPRI research considered the following information sources to determine whether there is an existing list of vulnerabilities that could be used as a common set of vulnerabilities for attack surface analysis that are consistent with the first principles:

- NIST Special Publication 800-82, “Guide to Industrial Control Systems (ICS) Security,” Revision 2, May 2015 [20]
- Common Weakness Enumeration, <https://cwe.mitre.org/data/index.html> [21]
- Common Vulnerabilities and Exposures (CVE), <http://cve.mitre.org/about/terminology.html> [22]
- Common Vulnerability Scoring System, <https://www.first.org/cvss/specification-document> [23]
- IEC 62443 series [10]
- National Vulnerability Database (NVD), <https://nvd.nist.gov/cwe.cfm> [24]
- Open Web Application Security Project, <https://www.owasp.org/index.php/Category:Vulnerability> [25]

Most of the vulnerabilities from the sources above are described as a software weakness or flaw, or lack of implementing a cyber security control. A specific software weakness or vulnerability (such as those published in the NVD) may be exploited by an adversary using some form of malware; however, the specific software weakness does not in itself cause harm. EPRI found that the characteristics of existing vulnerability schemes described in these common references make them unsuitable for use in technical assessments of digital assets. In fact, the universe of existing vulnerability schemes, including known and unknown software vulnerabilities, is unbounded. This situation is aptly described in an article by Pascal Meunier, titled, *Classes of Vulnerabilities and Attacks* [26].

No matter the type of vulnerability or exploit, an adversary is attempting to achieve certain objectives. “The Industrial Control System Cyber Kill Chain,” Michael J. Assante and Robert M. Lee, October 2015, states that an attacker’s objectives fall into three categories: loss, denial, and manipulation [27]. “They include a loss of view, denial of view, manipulation of view, denial of control, loss of control, manipulation of control, activation of safety, denial of safety, manipulation of safety and manipulation of sensors and instruments” [27]. From a more systematic perspective, an adversary attempts to exploit certain properties of an asset to cause harm through theft or alteration of digital information that includes changing asset operation/behavior, presenting false information, or by directly affecting asset operation. This is an adversary’s exploit objective.

In order to achieve the objective, an adversary seeks to identify and understand to a deep technical level the existing systems, systems configuration, communications, and functions in order to launch a successful attack. The adversary is using a systems engineering approach to understand and compromise the facility. Consequently, defenders must understand their critical systems using a systems engineering approach as well as or better than adversaries to provide effective and efficient protection.

A systems engineering approach that fully develops the system composition, system decomposition, system and device configuration, and data flow allows an engineer to know the attack pathways and exploit mechanisms that could be used by an attacker to achieve the objective. The result of this analysis is the attack surface. An adversary must use an existing attack pathway and exploit mechanism in order to achieve the objective. If all attack pathways and exploit mechanisms are mitigated to an acceptable risk level, then the attack surface is mitigated regardless of any software vulnerabilities.

Example 4-1

Attack surface analysis to mitigate a known software flaw

Consider CVE-2014-9191 that affects numerous HART transmitters that are commonly used in fossil power plants. “The CodeWrights HART Device Type Manager (DTM) library in [vendor] HART DTM [device] allows physically proximate attackers to cause a denial of service (DTM outage and FDT Frame application hang) by transmitting crafted response packets on the 4-20 mA current loop.”

- *A software flaw in the DTM library could allow an adversary to interfere with the normal operation of the asset by presenting false demands for asset interaction; known as a Denial of Service (DOS) attack.*
- *What is interesting is that in order to “transmit crafted response packets,” an adversary must use an existing attack pathway and exploit mechanism.*
- *Attack surface analysis shows that normal communications with a HART transmitter requires a physical connection to the analog loop (physically proximate) to transmit HART data packets (response packets) using a HART configuration device.*
- *An exploit of the DTM vulnerability (HART configuration device and software) must use the loop physical access (Attack Pathway) and HART protocol (Exploit Mechanism) to initiate a DOS attack, regardless of how clever the design of the exploit.*
- *Controlling physical access to the analog loop and unauthorized communication using the HART protocol mitigates the DTM software vulnerability by mitigating the attack pathway and exploit mechanism.*

An attack surface is characterized by using a systems engineering approach to identify the attack pathways and exploit mechanisms. Owners should mitigate the existing attack pathways and exploit mechanisms that allow an adversary to achieve an objective, rather than postulating and protecting against all potential “Exploits” or “Vulnerabilities.”

Both CCE and TAM identify an attack surface using a systems engineering approach. The Sandia Laboratories IDART assessment process also uses a systems engineering approach. Their approaches to identifying the attack surface vulnerabilities, however, vary.

CCE Attack Surface Identification Best Practices

The CCE methodology uses the outputs from Phases 2 and 3 in concert to identify the attack surface, attack vectors, and their associated vulnerabilities. Phase 2 examines the system from the organization’s perspective; Phase 3 examines the same system from the adversary’s perspective. This approach allows asset owners to understand what they do and do not know about their system, identify what critical information is protected and what is public facing, and discover choke points (or key attack vectors) that the adversary must traverse to accomplish the HCEs identified in Phase 1.

Phase 2, or “System-of-Systems Analysis,” is primarily a data collection effort on your own system. Phase 2 examines the specific system(s) involved in each HCE identified in Phase 1 through a systematic analysis of all related equipment, systems, processes, operations, maintenance, testing, and procurement practices [28]. The team collects, organizes, reviews, and summarizes the information related to every aspect of the system(s). Recall that Phase 1 already prioritized the most critical components of the business, so the data collection is focused only on systems relevant to the adversary’s goals.

Phase 2 requires that asset owners fully understand their system(s), and asking the right questions is usually the best place to start. If the HCE requires access or manipulation of a specific human-machine interface (HMI), the first questions should answer the standard *who*, *what*, *where*, *why*, and *how* of the system. Who has access to the HMI? What kind of access? Where is the HMI physically located? Where is the HMI within the system? Geographically? Logically? What network is the HMI on? Who has access to that network? How are the supply chain (during pre-commissioning) and replacement (during mid-life) processes handled? What version of firmware is the HMI? Does the vendor have access to the HMI for updates/servicing? When is the maintenance schedule? Who has access to the physical location of the HMI? How often? Are operators and maintenance workers using shared credentials?

The questioning does not stop there. The goal is to fully understand everything that the organization knows about its processes and procedures. This entire phase is heavily driven by Systems Engineering Principle 4 (**SE4**) at the beginning of this section. To reiterate, organizations must understand their own systems. The organization must fully understand its systems in order to effectively characterize and mitigate the attack surface.

In contrast, during Phase 3, “Consequence-based Targeting,” the team develops plausible attack scenarios by examining the HCE and information collected in Phase 2. This time, however, the team is looking at the key systems and processes from the adversary’s perspective. The attack scenarios are fully developed by the CCE Team using Phase 2 information to identify what specific choke points the adversary would need to cross to cause the HCE.

To fully develop these attack scenarios and identify choke points, the team uses several tools developed in the CCE methodology as strategies to identifying, visualizing, and ultimately mitigating the attack surface. One key piece for visualizing and understanding the specific steps that an adversary must accomplish to cause the HCE is the CCE kill chain. Assembling the cumulative knowledge, capability, and access that are needed to maliciously manipulate a system requires a long-term “campaign” of iterative targeting and information collection activities [28]. The CCE kill chain (Figure 4-3) helps the team visualize the adversary’s critical requirements.




Figure 4-3
The CCE kill chain describing Phase 1 through Phase 3

The CCE kill chain helps illustrate the specific information requirements the adversary needs to develop and deploy a payload to cause an HCE. By developing attack scenarios, the CCE Team can begin to understand the specific targets the adversary needs to access. These targets can be components, processes, memory modules, programmable chips, or logic circuits. Also included in a list of potential targets are the non-cyber parts of the process, like personnel with direct access to the system [28].

The CCE methodology further breaks down these targets by defining technical approaches to each target. This helps understand the attack surface and vectors by describing the access, actions, timing, and triggering that the adversary requires to manipulate the target.

The goal of Phase 3 is not to eliminate every one of the adversary's entry points into the system. This is not practical (or possible!). Recall Risk Management Principles 3 and 4 (**RM3**, **RM4**). It may not be possible to eliminate cyber risk, but the impact and likelihood of the attack can be mitigated by discovering and implementing specific controls at strategic locations in the adversary's attack scenario. This is done in CCE later in Phase 4, but by accepting that we cannot close every door, asset owners take the first step in the effective, affordable, and calculated cyber security by focusing efforts on what specific choke points will obstruct the adversary's ability to cause the HCE.

Risk Management Principle 2 (**RM2**) in OT cyber security also warrants a mention here. Cyber risk changes over time. As mentioned previously, CCE is designed to be iterative. Even the most thorough system-of-systems analysis may overlook a component of the process. By performing CCE iteratively through each stage of the generation plant's life cycle, plant operators can continually monitor and record key choke points for adversaries. Recall that identifying these choke points allows the plant owner to apply mitigations, protections, and controls where they are needed.

TAM Attack Surface Identification Best Practices

Step 1 of the TAM requires the user to characterize the attack surface and identify exploit sequences using a systems engineering approach. The attack surface of a device or system is defined by identification of the existing exploit sequences within a defined scope. The attack surface is documented as Part 1 of the CSDS.

The TAM begins with a close examination of the asset composition in its installed configuration, including its hardening features, and how data actually flow in and out. The installed configuration and data flow then bound the assessment and the applied security control methods. Technical information about the asset is required, but the TAM can make use of the best information available. The assessment scope depends on the asset composition, asset decomposition, the available technical information, and the installed configuration and data flow, as illustrated in Figure 4-4.




Figure 4-4
Assessment scope

After the asset composition is identified, the TAM characterizes its attack surface and communication relationship with other assets in order to identify attack pathways. Attack pathways are used by attackers to (1) take direct action against an asset or (2) steal or alter *critical data*, such as operational data, configuration data, or security data. Critical data can be at rest in an asset or in transit between assets. The goal of an attacker is to take direct action or steal or alter critical data to achieve an exploit objective. When attack pathways are identified, the mechanisms used to exploit those pathways (that is, achieve an exploit objective) are identified. Taken together, each unique combination of exploit objective, attack pathway, and exploit mechanism forms a distinct exploit sequence:

$$\text{Exploit Sequence} = \text{Exploit Objective} + \text{Attack Pathway} + \text{Exploit Mechanism}$$

Discovering and documenting the installed configuration and data flow is one of the most important activities of the TAM. A properly documented data flow bounds the scope of the analysis and creates significant efficiencies in implementing the rest of the TAM. Without a fully documented and bounded data flow, an assessor may end up analyzing multiple configurations, use cases, and communications that are not installed or in use, or become confused about what to analyze.

The installed configuration and data flow comprise the following three areas:

- Installed configuration
- Critical data at-rest and in-transit
- Data topology and data flow

The installed configuration and data flow are a key part of bounding the assessment scope in Step 1. Developing the installed configuration and data flow occurs in parallel with asset composition and decomposition and all other aspects of the TAM. Five activities are performed to identify the installed configuration and data flow, as shown in Figure 4-5. EPRI has published a separate *Cyber Security Data Flow Identification and Documentation Topical Guide* (3002015761) [29] that provides more detailed guidance on how to develop an installed configuration and data flow.




Figure 4-5
Activities to identify the installed configuration and data flow

A key activity in Step 3 of the TAM is performed in conjunction with identifying the installed configuration and data flow in Step 1, identifying *Connectivity Relationship Sets*. An asset is a member of one or more connectivity relationship sets that help identify communication pathways and components in each pathway that an attacker could use to achieve an exploit objective. Developing connectivity relationship sets coincidentally with the data flow is highly efficient and allows the owner to visualize the installed configuration and data flow.

Asset features and data flow are needed to characterize the attack surface, including features that support manual data entry or retrieval via menus on a display or pushbuttons on a faceplate. When characterizing the attack surface, the assessor will identify attack pathways that could be exploited if they are not secured.

An exploit sequence is an attack pathway and exploit mechanism that allows an attacker to achieve an exploit objective:

$$\text{Exploit Sequence} = \text{Exploit Objective} + \text{Attack Pathway} + \text{Exploit Mechanism}$$

An adversary will attempt to use one of the available attack pathways and existing device or system capabilities to achieve an exploit objective. Attack pathways are enabled via the attack surface characteristics identified earlier. Attack pathways are used by attackers to (1) take direct action against an asset or (2) steal or alter critical data, such as operational data, configuration data, or security data. Critical data can be at rest in an asset or in transit between assets. The particular asset capability that allows an attacker to achieve an exploit objective is defined as an exploit mechanism. Each unique combination of exploit objective, attack pathway, and exploit mechanism is an exploit sequence.

Unmitigated exploit sequences can lead to consequences if they are exploited, and their identification allows for a systematic and more specific assessment of devices and systems, as well as structured allocation of security control methods with a consistent technical basis.

An exploit objective can be a direct action on the device itself, or it can be theft or alteration of critical data. The goal of an attacker is to take direct action or steal or alter critical data to achieve an exploit objective.

IDART Attack Surface Identification Best Practices

The Sandia National Laboratories IDART [30] also uses a systems engineering approach to identify an attack surface. The IDART method was developed at Sandia National Laboratories and is a NIST-recognized method in NIST SP 800-115, “Technical Guide to Information Security Testing” [31]. IDART is primarily focused on assessment of security, enterprise, and control systems, and so on, to understand how adversaries might target the system and accomplish their goals through identification of attack pathways and exploit mechanisms.

IDART assessments include the following:

- Characterizing the target system and its architecture
- Identifying nightmare consequences
- Analyzing the system for its security strengths and weaknesses
- Identifying potential vulnerabilities whose exploitation will result in nightmare consequences
- Providing prioritized mitigation strategies so owners can make risk-informed choices

IDART provides a tool for understanding the system being evaluated and consequences of concern.

Best Practices for Risk-Informed Attack Surface Mitigation

This section describes a cybersecurity strategy that is based on implementing the most cost-effective and sustainable security control methods to mitigate an attack surface to an acceptable risk level. Best practice attack surface mitigation uses the principles of the NIST Framework [32] to identify, protect, detect, respond, and recover to an acceptable risk level and cost for the business.

Identify. The collective set of information gathered as part of the systems and risk analysis and is not directly related to a control method security function. Security control methods are implemented for three security functions, as follows:

- Protect
- Detect
- Respond and recover

Protect. The protect function is used to evaluate the effectiveness of security control methods that protect against unauthorized physical and/or logical access to a device or system, and unauthorized changes or activity. The PIN code entered for access to a protective relay is an example of an engineered security control method intended to protect against unauthorized access to the protective relay.

Detect. The detect function is used to evaluate the effectiveness of security control methods that support the detection of anomalies and/or a compromise that may be an indication of an unauthorized activity. Some events captured by digital components can provide information useful in detecting possible cyber security incidents. Consider a configuration change event log in a transmitter that captures the last time that the configuration was changed. This event log may provide useful information for the detection of unauthorized access and/or modification, although after the fact, if the transmitter is not continuously connected to an Engineering Tool Workstation.

Respond and recover. The respond and recover function is used to evaluate the effectiveness of security control methods that support response activities, such as event log forensics, and recovery activities, such as restoring a configuration baseline. For example, if during recovery it is discovered that the event records indicate that several configuration change events have occurred over the last six weeks, then incident investigation can focus on activities that may have resulted in the event, and recovery efforts may focus on restoring from a configuration backup performed prior to six weeks ago.

From an owner's perspective, there is a difference between cyber security requirements and cyber security methods. Several cyber security regulations, standards, and guidance are based on applying a catalog of cyber security control requirements, but do not provide guidance on how to implement the requirements. Different devices and systems have different capabilities that may or may not be used to meet a requirement, and some requirements may not apply for a given component or system.

There is a significant challenge to interpret a cyber security control *requirement* into a meaningful *security control method*. Best practices identify the specific attack pathways and exploit mechanisms, and then consider the available device and system features and functions, as well as site-level cyber security control methods to mitigate the attack pathways and exploit mechanisms, rather than start with a set of cyber security control requirements. In addition, the cost to implement and maintain a control method must be considered.

Example 4-2

Control requirement does not provide guidance on how to implement a control method.

Consider the NIST Framework Data Security Category: “PR.DS-1: Data-at-rest is protected.” This is a requirement that can be implemented in many different ways. For a Windows workstation running the SQL Server, the workstation disk can be encrypted via “Device Encryption” and data in a column in a SQL Server database can be encrypted. For a pressure transmitter, there is no encryption function nor can an outside encryption tool be used; therefore, other control methods must be used such as restricting physical access, configuration control, and configuration change event logs to protect data-at-rest in the transmitter.

Requirements do not provide guidance on how to meet the requirement with specific control methods. Owners should first identify the attack surface and evaluate and allocate the available control methods to mitigate the attack surface. After the attack surface is mitigated, it is straightforward to demonstrate how the mitigation meets the intent of the requirement.

Example 4-3

Control requirement cannot be implemented or may not apply.

Next, consider the NIST Framework Security Continuous Monitoring Category: “DE.CM-8: Vulnerability Scans are Performed.” For many ICS devices, a vulnerability scan is not possible (for example, pressure transmitter running with proprietary firmware) or an active scan will interfere with the device operation (for example, anti-malware scan of a PLC in operation results in plant equipment malfunction.). What is the purpose of the scan, and how should the scan be performed? Active or passive scan? Anti-malware or unpatched vulnerabilities? Abnormal behavior?

A cyber security strategy that focuses only on security will attempt to implement all available control methods. However, the purpose of cyber security is to protect the business, workers, and the public, and should be balanced against successful sustainable operations. The cost of installing and maintaining a security control method can have a large influence on whether to implement a security control method or consider another security control method to achieve the same outcome.

Even if the security effectiveness of a security control method is high, the cost to implement the security control method may be high enough to consider implementing a different security control method. In contrast, if the security effectiveness is low and the cost is also low, then it may be worth considering implementing the security control method even though it has a low security effectiveness. Initial and ongoing costs need to be considered. Systems engineering best practice allows an owner to choose between technical options to balance security effectiveness and cost, to achieve an acceptable business outcome.

The following provides descriptions on the effectiveness of the control methods and how an owner knows when the attack surface is mitigated to an acceptable risk level:

- Control methods and the combination of control methods are qualitatively scored for security effectiveness through a formal approach or by expert consensus. For example, a four-digit PIN for access control is less effective than an account with least privilege and a strong password. Combinations of control methods have a higher security of effectiveness than any single control method such as a locked cabinet inside a locked room with a password to access a PLC in the cabinet.
- Risk is factored into an allocation of control methods to determine the score needed to reach an acceptable risk level. For example, the consequence of theft of cooling water intake temperature values is much less than theft of the plant control system network architecture. The control method score required to meet an acceptable risk level for theft of meteorological data is less than theft of the network architecture.
- Both the plant systems and control methods are temporal in nature and must be accounted for in control method scoring and risk. Consider a Windows XP laptop that is required for maintenance of an older device. Windows XP is no longer supported with security patches, and it is easy for an adversary to craft an exploit; therefore, the exploit difficulty is lower, which increases the risk. However, if the device has been designated as a run-to-failure, then the consequence is lower, which lowers the risk associated with a Windows XP exploit.
- When all identified attack pathways and exploit mechanisms are mitigated to an acceptable risk level, the owner can be confident that the device, system, and plant are secure.

A cyber security strategy that allocates cyber security control methods specific to the attack surface and that considers and adapts for system, facility, and component life cycles will be more flexible, reduce costs, and provide greater mitigation than a single rigid software suite or penetration test.

CCE Risk-Informed Attack Surface Mitigation Best Practices

Phase 4 of the CCE methodology, titled, “Mitigations and Protections,” guides generation plant owners through developing mitigation and protections to stop, thwart, or stall the adversary. Recall that Phase 1 identified HCEs that represent the consequences of cyber-enabled sabotage that matter most to the asset owner. Phases 2 and 3 directed efforts at developing a knowledge base and a thorough understanding of the system from both the insider’s and the adversary’s perspectives.

This phased approach leaves asset owners and the CCE Team with a prioritized list of the consequences of cyber-enabled sabotage, a comprehensive understanding of the systems involved in the sabotage, a collection of what the adversary can find out about those systems in open source, and fully developed attack scenarios describing what an adversary needs to accomplish. As a result, specific choke points in the system that the adversary *must traverse* (either logically or physically) are identified.

Now it is time to allocate resources for triage. Fortunately, the list of critical choke points is already prioritized based on their individual impact to the organization—the fruit of the consequence prioritization labor in Phase 1. The chosen mitigation mechanisms should be developed to either enhance the organization’s ability to detect adversarial actions, increase the adversary’s costs, or decrease the organization’s cost in the event of a successful attack.

Phase 4 in CCE borrows heavily from the NIST Framework [33]. CCE mitigations and protections are categorized by function [34]. Since Phases 1 through 3 cover the “Identify” function, the remaining functions describe the categories of mitigations and protections available to the CCE Team.

To better address CCE’s goal of protecting critical infrastructure, CCE uses its own definitions for these four functions; however, the principle behind each remains largely the same [34]:

1. **Protect.** The ability to remove the objective of cyber-enabled sabotage (take it “off the table” for an adversary).
2. **Detect.** Enables timely discovery of adversary activities.
3. **Respond.** The ability to contain or disrupt adversary activities.
4. **Recover.** Timely restoration of critical functions and services.

The protect function is the most valuable to the generation facility. If there is an affordable, practical protect option on the table, it should be implemented over any other mitigation effort. If used correctly, it will remove the adversary’s ability to cause the HCE through cyber means.

Example 4-4
Protect function [34]

Consider a liquid storage tank at an industrial facility. A PLC governs the pump controlling the fill level of the tank. If an overflow of this tank is a catastrophic HCE, a possible protect mechanism is the use of a separate float switch in the tank. This switch, when activated, will physically disconnect the pump from the power supply. By installing such a device, an adversary will not be able to cause this HCE by cyber means alone.

The detect function is secondary to protect mechanisms. Detect options are designed to notify operators and owners of adversary activity early on—before damage reaches irreparable thresholds, that is, before the adversary can cause an HCE. A well-designed detect trigger will give the good guys a chance at stopping the HCE when a protect mechanism is not practical or affordable.

Example 4-5 **Detect function [34]**

An excellent, real example of effective detection measures can be found in Ukraine's response to the December 2015 cyber attack. A Ukrainian power company detected adversarial actions in the corporate network. They were able to act, which may have spared that specific organization from an attack in the end. Of the actual victims of the attack, not one detected adversary activity before the actual power outages.

The respond and recover functions of the CCE methodology speak more to the organization's readiness and preparedness for an attack. A response plan will help everyone in the organization know what to do during an attack, which will help contain or disrupt the adversary's actions during the attack. In contrast, a recovery plan outlines responsibilities and operations that the organization will perform after an attack. The purpose of a recovery plan is to return to normal—or at least acceptable—operations as soon as possible.

There is no all-purpose method for prioritizing mitigations and protections. Each organization's systems, organization, needs, and capabilities vary; however, as a best practice, consider the following [34]:

- **Type (“Protection” vs. “Mitigation”).** Protections will prevent an adversary from causing an HCE via cyber means. Mitigations cannot do so.
- **Efficacy.** Proposed protections or mitigations should be reviewed for their perceived efficacy—whether the proposed solution makes the attack not feasible, or to what degree it can reduce any negative consequences or make things more costly or challenging for an adversary. This kind of review is limited in that the efficacy of a solution ultimately relies on the specific implementation that is adopted.
- **Existing threat information.** Some attack scenarios may leverage techniques that have already been witnessed “in the wild” (deployed against a victim) or involve targeting of systems or components that correspond with known adversary interest. The presence of existing capability or research directed against these systems will presumably result in an increased likelihood that an adversary would first pursue these options over other scenarios.
- **Assessed attack difficulty.** An increase in attack difficulty will presumably correspond with a decrease in the likelihood that an adversary would first pursue such a path, due to the increase in relative cost and/or the increased need for specialized skills or knowledge. In such an evaluation, the difficulty ranking can be baselined at the level of the least challenging scenario.

CCE is intended to serve as a triage activity. In some cases, it may not be possible to prevent an adversary from achieving their objective. In such cases, defensive actions should focus on increasing the resources required of an adversary to perform cyber attacks or decreasing the resources that a victim organization will need for recovery purposes [34].

TAM Risk-Informed Attack Surface Mitigation Best Practices

In the TAM Step 2 and Step 3, exploit sequences identified in Step 1 are mitigated using selected, available security control methods. Allocation of security control methods is informed by risk through exploit difficulty, security effectiveness, and facility consequences. An owner allocates security control methods using implementation burden as a cost indicator to select from competing alternatives.

Control Method Security Effectiveness

Individual security control methods will have more or less security effectiveness for mitigation of exploit sequences. *Security Effectiveness* of a control method is a function of its implementation effectiveness (that is, for each security function, how effective is the control method if it were to be implemented?) and the exploit difficulty it imposes (that is, how hard is it for an attacker to overcome the control method?). Exploit difficulty is the inverse of the “likelihood” that a control method can be overcome. Using exploit difficulty as a surrogate for “likelihood” is a component of creating a risk informed allocation of security control methods:

$$\text{Security Effectiveness} = f(\text{Implementation Effectiveness}, \text{Exploit Difficulty})$$

Security effectiveness is scored using the guidance, criteria, and formula provided in the TAM. Owners allocate security control methods to exploit sequences that comprise the attack surface. But what control method security effectiveness score is needed in order to determine whether the attack surface is mitigated, and is more than one security control method needed?

Combined Security Effectiveness

A *Combined Security Effectiveness Score* is calculated for one or more control methods allocated to an exploit sequence. As attackers move along an attack pathway, they will encounter a series of obstacles in the form of control methods. An entropy-based formula is used to calculate a score for the combination of control methods allocated to each exploit sequence. A benefit of using an entropy-based score is that entropy is, by definition, an additive so that, as each control method is encountered, the entropy associated with exploiting each control method is simply accumulated. Combined security effectiveness is scored using the guidance, criteria, and formula provided in the TAM.

The individual security effectiveness scores do not consider the possibility that, after a control method is defeated, further implementations of control methods along an attack pathway become increasingly less effective. However, as additional implementations of a control method occur, the difficulty of defeating that control method decreases.

The combined security effectiveness score is different than an individual control method security effectiveness score in that it provides a geometrically weighted sequence that allows for conservatism when combining multiple control methods for an exploit sequence.

Risk-Informed Target Level Threshold

A combined security effectiveness *target level threshold* is determined by a consequence analysis method chosen by the facility owner. There are five default target levels A–E, with A representing the highest consequence target level. Nominal threshold values for combined security effectiveness are provided for each target level. The number of target levels and associated combined security effectiveness threshold values can be changed by the owner, based on the type of hazard analysis used and the owner’s risk tolerance.

The combined security effectiveness score includes exploit difficulty, which is the inverse of likelihood in the risk calculation; therefore, a higher combined security effectiveness score increases the exploit difficulty, thus reducing the likelihood of a successful attack. Using this risk-informed method allows an owner to know when the attack surface is mitigated for a given consequence while only allocating the most effective security control methods, resulting in an efficient and cost-effective process.

An additional benefit of using the target level threshold is that a higher threshold corresponds to an adversary’s capability. In other words, a higher target level threshold requires a higher adversary threat capability.

An owner with a lower risk tolerance may choose to raise the threshold values for one or more target levels. An owner with a higher risk tolerance may choose to lower the threshold values for one or more target levels. An owner may change the number of target levels and combined security effectiveness threshold values based on the facility owner’s hazard analysis, owner requirements, and risk tolerance.

EPRI has published a separate *Risk Informed Target Level Topical Guide* (3002015760) [7] that provides more detailed guidance on how to use consequence analysis and target level thresholds in the TAM.

Control Method Efficacy

When multiple security control methods are available and various combinations of methods meet or exceed the combined security effectiveness target level, the owner can use control method efficacy scores to select from competing alternatives. The *efficacy score* of a security control method or combination of methods is a function of its security effectiveness score and implementation burden:

$$\text{Security Control Method Efficacy Score} = f(\text{Security Effectiveness}, \text{Implementation Burden})$$

The *implementation burden* is an indicator of the cost to implement a control method. The security effectiveness score of each control method or combination of methods has already been determined at this point, but to determine its efficacy, the owner must also score its implementation burden using the guidance, criteria, and formula provided in the TAM.

The *implementation burden* is scored based on the *initial implementation burden* and *ongoing implementation burden* or whether implementation of a control method *conflicts* with plant operation. When the implementation burden for a security control method is scored, the owner can then determine the control method efficacy score on a scale of 1 to 5, with 5 being the

highest efficacy. Note that the efficacy score of each security control method is determined three times, once for each of the security functions (protect, detect, respond, and recover), which will likely result in a different efficacy score for each function.

Table 4-1 provides a “heat map” that correlates the security effectiveness score of an identified control method with its implementation burden score to produce an overall efficacy score.

Table 4-1
Security control method efficacy scoring

Control Method Efficacy Score (Three Times, Once for Each Security Function)				
Security Effectiveness Score	Implementation Burden			Conflict
	High (>2.3 to 3.0)	Medium (>1.3 to 2.3)	Low (0.0 to 1.3)	
None	None			Do not Implement
Low (0.1 to 1.0)	1	2	3	
Medium (>1.0 to 2.0)	2	3	4	
High (>2.0 to 3.0)	3	4	5	

Efficacy scores are determined at the device level and allow an owner to select from competing alternatives and begin allocating the preferred security control methods.

Allocate Preferred Engineered Security Control Methods

Allocation means deciding which security control methods to credit for mitigating each of the exploit sequences identified in Step 1. In Step 2, security control methods that are features and functions of a device that can be allocated at the device level for each exploit sequence are termed *engineered security control methods*. If the allocations at the device level do not yield a security effectiveness score that meets or exceeds the combined security effectiveness target threshold value, the result is a *residual exploit sequence* that is subject to further treatment via shared security control methods in Step 3 of the TAM.

Shared Site-Level Security Control Methods and Relationship Sets

When the allocated engineered security control methods don't meet the target level threshold for any identified exploit sequence and security function, then that exploit sequence is residual and must be mitigated via site-level security control methods.

Shared control methods are owner implemented site-level control methods (for example, configuration management and transient cyber asset programs) that are not included as features and functions of the device. Shared control methods may include physical protection, such as locked and alarmed rooms or cabinets, or administrative procedures, such as patch management or supply chain procedures.

The owner determines the shared control methods to develop or apply and come from a predetermined set of control methods that can be referred to as a *shared control method library* (Shared CML). Shared control methods are inherited by a device or system through a relationship set. *Residual exploit sequences* are mitigated using site shared control methods.

A device can be deployed in a facility either in a standalone configuration or as a component in a collection of components in a *relationship set*. A relationship set allows for easier analysis and optimizing allocation of control methods through *inheritance*. A device can be a member of more than one relationship set.

The benefit of creating a relationship set allows a vendor or owner to:

- Efficiently and cost-effectively scale the TAM across systems and the facility.
- Group devices by relationship criteria to ensure that the common kill chain is mitigated, including connectivity, spatial, and programmatic exploits through inheritance of control methods.
- Allow an owner to increase the security margin of system or component if desired.
- Identify *normalized exploit mechanisms* to reduce the proliferation of similar but different exploit mechanisms to improve the efficiency of analysis.

All devices inherit the exploit sequences and control methods (both local and shared) of all other assets in a relationship set through aggregation. This key aspect of the TAM ensures that all relevant kill chains are mitigated without excessive expansion and redundancy in the CSDS.

IDART Risk-Informed Attack Surface Mitigation Best Practices

IDART [30] uses a team of SMEs to evaluate and allocate potential mitigations for an identified attack surface. Using an expert panel approach, the team prioritizes control method mitigation strategies based on control method effectiveness and system or component risk.

CCE Best Practice Case Study

Typically, organizations that implement the CCE methodology prefer to not disclose their systems, vulnerabilities, mitigations, and operating procedures; moreover, they would also prefer to not disclose the functions and services they deem most critical to their operations.

This hesitation to share what worked (and what did not) with others in industry is further delayed when dealing with organizations within the nation's critical infrastructure, where concerns over dollar signs shift to those accompanied with more grave consequences to safety and national security.

In response to this conundrum, INL has drafted several fictional case studies loosely related to real organizations that have implemented CCE methodology successfully both on their own and with the assistance of INL. The organization-specific details are obfuscated but largely irrelevant to gaining an understanding of the CCE methodology. What does matter are the lessons learned—those lessons that are largely applicable across nearly every critical infrastructure

sector. Whether these lessons point to a specific success or how an organization overcame a hurdle, the payoff is the same. CCE is a triaged approach that reinforces the use of cost-effective mitigation strategies where they matter—where the asset owners have deemed most important to their specific needs.

DISCLAIMER #1

This case study is a work of fiction. It is the product of the authors' imaginations, written to reinforce the understanding of the CCE methodology. Names, locations, events, corporations, regions, countries, and incidents are fictitious. Any resemblance to actual countries or events is purely coincidental.

Baltavia Power Outage [30]

Fictional country **Baltavia** is undergoing the lengthy process of modernizing its transmission substations with new auxiliary dc power management systems. While the systems will offer enhanced capabilities, like remote control and remote monitoring capabilities, Baltavian officials are more concerned with the power management systems' increased reliability.

Demonstrating competence to its European neighbors is its primary concern. Baltavia seeks to establish a footing as a key energy provider to European markets. Competing nation-state energy providers in the region understand this goal and may pursue the inverse—a campaign using cyber-enabled sabotage to cause a substantial outage. This outage would damage Baltavia's credibility, as well as likely ruin their market prospect altogether.

It is not enough, nor is it often possible, to designate efforts to defending from “all” cyber-outages. Even if it was obtainable and the organization had unlimited resources, recall the evolving nature of cyber attacks discussed in the first principles. By the time years and millions have been spent securing the system from the threat, the threat—as well as the adversary's approach—would have evolved. This is not costly nor practical for organizations.

Baltavia realizes that, while power delivery is still possible with an outage at a single key delivery point, an outage exceeding 6 hours would be catastrophic, impacting throughput and system resilience capabilities. This would almost certainly erode the EU's perception of energy security in Baltavia and damage their perception of Baltavian energy security. A shorter outage of 2 hours or less may be recoverable through transactions by government officials and utility management [30]. As such, we arrive to a specific boundary condition that will drive future CCE efforts:

Functional disruption at a single critical transmission substation, resulting in a reduction of Baltavian full power delivery capabilities to Eurozone markets for 6+ hours [30]

In this case study, the boundary condition above was used to draft dozens of events and an equal number of cyber-events—that is, events that are possible through cyber means. Essentially, the CCE Team was able to discover several different ways an adversary *could* affect Baltavia's system in a way deemed disastrous to the nation's success. After a full round of Phase 1 Consequence Prioritization, one cyber-event became the sole HCE because its consequences were the direst to Baltavian officials.

HCE: The adversary gains access to the substation network and manipulates configuration of the battery management system. Modifications reduce battery bank re-charging capability as well as dc power availability. The dc system capacity is degraded to a level insufficient for sustained support of substation SCADA, protection, and operations infrastructure. Attack ensures that no indications are presented to system operators while the charging system is at reduced capacity. Adversary triggers station isolation and de-energization by opening breakers on three transmission-level interconnects. [30]

Phases 2 dissects the HCE identified in Table 4-2 from Baltavia's perspective. Documents are gathered on the physical system, the control system design, the software/firmware, automation/control center systems, remote management systems, wiring diagrams, network architecture, and personnel involved with O&M. Fully developed system descriptions are drafted for the transmission system, the critical substations, system operations, system communications, substation control operations, and the substations' new auxiliary dc power systems.

Table 4-2
Severity scoring matrix of the Baltavian case study HCE (cyber-event #4) [30]

		Severity Scoring			
		None (0)	Low (1)	Medium (3)	High (5)
	Attack Breadth $\beta = 1$		Elements of the system are vulnerable to an exploit that is active and causing operational effects, but recovery is possible using immediately available resources. These events are covered within the utility's recovery plan.	Multiple system elements have the potential to be or have been successfully attacked causing operational effects. Recovery is possible but requires additional resources (that is, time, personnel, and so on) not immediately available.	Many system elements have been successfully attacked causing operational effects. Restoration is complicated by the dispersed deployment of devices or scale. Timeline for recovery is unknown.
	Duration $\delta = 3$		Return of all service in less than 2 hours.	Return to service in between 2 and 6 hours.	Return to service in greater than or equal to 6 hours.
	System Integrity Confidence $\epsilon = 2$		Asset owner has ability to restore and is confident in restoration integrity.	Asset owner has knowledge to restore but does not have the resources (financial, time, personnel, and so on) to restore confidence in the system.	Asset owner has ability to restore but is not confident in restoration integrity.

Phase 3 approaches the HCE and the systems involved from the adversary's perspective. To understand the adversary's critical needs for deployment and development, Baltavia develops attack scenarios that could plausibly result in the HCE—a success for the adversary and a critical, mission-ending failure for Baltavia. Using open-source information, reconnaissance, and social engineering, the adversary can discover nearly all the same information Baltavia could in Phase 2 with their insider advantage. The difference between what Baltavia has and what the adversary can find represents key information that should be further protected.

In Phase 4, Baltavia brainstorms protect, detect, respond, and recover options. Cost, effort, and time play key roles in prioritizing mitigation mechanisms. Baltavia moves forward with the following combination of engineering and operational controls as protect mechanisms:

Implement separate, non-networked device-level dc battery system monitoring that provides alarming for under/over voltage, as well as other critical dc system parameters. Validation occurs visually via local HMI and/or transducer display. Operator verifies that dc power system parameters are within (+/- X%) of desired ranges (and local control setpoint) on a fixed time basis by procedure (such as, "check setpoint hourly as part of rounds"). This would also require the operator to have a response procedure (detailed in section below). [30]

Additional detect measures are implemented as cost-effective methods of identifying an adversary's actions. New procedures are chosen for their cost-benefit ratio, such as device event monitoring, networking monitoring, improved endpoint malware detection, and account monitoring. Ultimately, because detect methods can be resource-intensive, Baltavia shifts to only those proven to be most effective in future operations.

Baltavia also develops strong incident response, management, and recovery plans as part of their respond and recover efforts. Baltavia chooses to incorporate best practices, like out-of-band communications infrastructure, and targeted training on bridging communication gaps between OT, IT, and corporate. In the event of the adversary's success, Baltavian officials and utilities have a clear roadmap of how to recover, including who makes what decisions and when.

TAM Best Practice Case Studies

Various generation plant owners have evaluated and implemented the TAM. Each owner was able to demonstrate benefit. Three TAM case studies are summarized below.

Rapid and Effective Cyber Incident Response Using the TAM

A generation plant held their first cyber security incident response drill, simulating an attack on a unit auxiliary transformer. The plant staff had previously developed a cyber assessment of the transformer according to the plant cyber security program that included the protective measures to be implemented for the transformer. The incident response drill was designed to test the plant's incident response capability, as well as test the completeness and effectiveness of the transformer protective measures. The incident response drill demonstrated the need to fully understand the system to a deep technical level.

The transformer assessment was based on high-level regulatory cyber requirements that included primarily administrative cyber security controls. From the beginning of the drill, it became clear that the existing transformer assessment did not contain the system information required to determine what happened during the attack. The incident response team had to investigate the

system by performing walkdowns and document reviews to discover the installed configuration and data flow. Without understanding the transformer configuration and data flow, the team could not determine the attacker's objective and how to respond. In addition, it was difficult to assess the effectiveness of the administrative and programmatic cyber security controls that were implemented.

Some months later during a TAM training class, the plant staff used the unit auxiliary transformer as the component evaluated during the TAM hands-on training exercise. At the end of the class, the cyber program manager and incident response team stated that had they developed the assessment using the TAM, all of the information needed to fully respond to the incident would have been at their fingertips. They estimated that the time to respond would have been significantly reduced and the effectiveness of the response would have been significantly higher. Furthermore, had the TAM been implemented, the cyber security control methods for protect, detect, and respond and recover tailored for the transformer would have demonstrated a high level of assurance that the transformer attack surface was mitigated.

TAM Cyber Engineering Approach Saved Over 1 Million Dollars in a Procurement

A generation plant had procured a new packaged system. Prior to installation, the plant staff wanted to assess the new system for cyber security compliance and asked the vendor to provide a proposal to assess and bring the new system into compliance. The vendor's proposal included a full analysis and custom modifications for over 1 million dollars. The high cost was unexpected. The plant staff decided to use the TAM to determine whether the vendor's modifications and services were needed and as a test case for the TAM.

The plant staff performed a full TAM analysis on the system as delivered. The analysis was completed in about one week. The attack surface was identified for the plant installed configuration and data flow. Existing system features and functions were evaluated and scored as engineered control methods. The site shared control methods that were applicable to the system were also identified and scored. The target level threshold was set as Level A, the highest level. Engineered and site control methods were allocated to the identified exploit sequences to determine whether the existing control methods were sufficient to mitigate the attack surface to Target Level A, and whether Target Level A demonstrated compliance.

The result was that the existing control methods were sufficient to meet the target level threshold and that Target Level A demonstrated compliance. The plant staff concluded that the custom modifications and services from the vendor were not required which saved over 1 million dollars and significant time in the schedule. The plant staff also concluded that the TAM approach was thorough and efficient, and proceeded to expand use of the TAM for the plant.

Thoroughness and Efficiency Using the TAM Cyber Engineering Approach

A utility owner needed to identify and mitigate the attack surface for a newly acquired DCS Security Suite that was being rolled out across the fossil generation fleet. The owner decided to evaluate the TAM through a pilot project for this complex system. The goal of the pilot was to evaluate the effectiveness, scalability, and usability of the TAM.

The project team attended the three-day TAM training course and implemented the TAM for the DCS Security Suite for a fossil generation plant. The security suite included multiple applications running in a virtual environment. The project team realized multiple benefits from implementing the TAM that included the following:

- The project team was able to more fully understand the system. This gave the team the confidence that they understood the attack surface as well or better than a potential adversary and that the attack surface was effectively mitigated.
- Unmitigated exploit sequences (vulnerabilities) were discovered and mitigated. The system had implemented mitigations; however, the entire attack surface was not fully mitigated.
- The Vulnerability and Patch Management Program was improved by better understanding software flaws (NVD/ICS CERT) and mitigation strategies as they arise. More importantly, the owner was able to reduce the number of patches that must be installed outside of scheduled maintenance outages, thereby reducing the cost and impact of the Patch Management Program.
- Security posture was improved while reducing cost by standardizing the DCS Security Suite configuration across multiple facilities. The CSDSs and RSDSs developed for the pilot project served as a reference assessment that is being reused across the fleet.

The utility owner decided to standardize use of the TAM because the TAM gave the owner higher confidence that the generation fleet would be more effectively protected against a cyber attack using a methodology that is more efficient and thorough than the previous process.

5

GAP ANALYSIS AND RESEARCH RECOMMENDATIONS

This section describes the technology, skill, and organization gaps that either hinder the implementation of best practices or are needed to enhance and accelerate cost-effective implementation of best practices. In developing the framework, the research team identified resources used for process (methodologies, standards, and guidelines) as well as tools and technologies used by OT cyber security practitioners in today's fossil generation fleet. Appendix A is a toolkit for practitioners to select the right resource for the need. The considerations outlined in Section 2, along with the Risk informed Process outlined in Sections 3 and 4, provide a basis for the practitioner to select appropriately.

As illustrated in Section 4, engineered (or technical) methods are more effective than administrative methods. For this reason, the toolkit includes 22 categories of tools and technologies, with commercially available and open-source examples used today. The list is not all-encompassing. The available capabilities are rapidly changing, which by itself can be a challenge for today's cyber practitioner. These technologies can enable a more effective and efficient implementation of the methodologies and standards.

With the toolkit providing the current state, several criteria for research gaps were identified. The criteria are discussed later in this section and were established through industry interviews and collaboration with end-users, leadership, technology providers, and research organizations. The noted gaps are areas of recommended coordinated research. The research team followed the process outlined in Figure 5-1.




Figure 5-1
Gap analysis process

Identifying Challenges and Research Gaps

Both the toolkit and gap analysis are divided into two primary domains, methods, and technologies. The first is methodologies, standards, and guidelines. This includes written guidance that fossil generation facilities use either directly (or indirectly) for programmatic or administrative controls. This can include high-level activities, such as developing policies or procedures, all the way down to specific application-level guidelines for performing cyber-related tasks.

The second domain is technologies and tools. These are hardware or software tools that physically, or via software or logic, support implementing and maintaining a cyber security program. This can include technical control measures for protection, detection, or response, and recovery. For example, the Intrusion Detection System (IDS) (available in many flavors) provides varying levels of detection capabilities and was considered a technology for this gap analysis. Other types of tools that were included in the review were software tools used to support specific processes. The publicly available Cyber Security Maturity Model (C2M2) software is an example of a tool that supports a programmatic maturity review, following the C2M2 methodology.

Standards, Methodologies, and Guides

Methodologies, standards, and guides serve as inputs or a basis for an OT cyber security program. Fossil generation facilities must choose those that align with their organizational goals, structure, business, as well as any required by regulation.

A **standard** is defined by Merriam Webster as “Something established by authority, custom, or general consent as a model or example” [35]. The NERC is an example of a body that has developed a set of standards, CIP Standards, required by regulation in the energy sector. When designing a cyber security program, mandatory and voluntary standards are considered as the minimal basis for the program.

One consideration with any standard is how the standard is being adopted by other stakeholders, such as vendors and technology providers. Many standards allow certification of products or services to the standard. While this provides some level of confidence, it is important to consider the specific field configuration of an asset. A certified, or regulatory compliant, digital asset may be configured in an insecure manner in the field. Neither certification nor compliance guarantees security.

A **methodology** is defined by Merriam Webster as “A body of methods, rules, and postulates employed by a discipline; a particular procedure or set of procedures” [35]. The TAM, as discussed in detail in Section 4, is an example of a methodology for assessing the cyber security posture and prioritizing mitigating control measures. It includes a procedure for performing the assessment, along with rules to guide the successful completion. In practice, these methodologies are adopted by an entity to support the cyber security program. As part of the implementation process, the entity tailors the procedure to fit their specific use cases.

A **guide** is defined by Merriam Webster as “Something that provides a person with guiding information” [35]. Many guides are available ranging from high-level, programmatic guides, to guides on specific equipment. These can be used in whole, or in part, with little to no modification for adoption. Vendors may provide guidelines on deployment and hardening associated with their equipment. Examples of standards, methodologies, and guides commonly used in the energy generation sector are included in Appendix A.

Tools and Technologies

While methodologies, standards, and guides dictate overall cyber-controls strategy, tools and technologies are used to physically, or logically, secure a system. Tools can include but are not limited to, software packages, as well as hardware. For this toolkit, tools are used to implement a methodology or are a cumulative group of technologies. Technologies on the other hand, are cyber mitigating controls. For example, performing identification by protocols, network characteristics, and physical MAC addresses can be portions of individual technologies, and then all these identification technologies may be grouped into a single tool.

When to Use Each

Methodologies should be used to enhance or create processes that solve a cyber issue. As an example, your organization may not have existing assessment processes in place. In this example, finding and adopting existing methodologies for assessment decreases the time required to create entirely new processes. Standards can either be required by a regulating body or be used voluntarily to improve cyber security readiness.

To continue the earlier example, new standards from a regulating body may have been released requiring your organization to develop assessment processes. Alternatively, after identifying that your organization is lacking assessment processes, standards, guides, or methodologies can be adopted and used as a starting point for creating assessment procedures based on existing best practices or regulations. To complete our example, once standards, guides, and methodologies have been identified and adopted, specific technologies can be implemented to provide mitigation against cyber attacks.

Challenge and Gap Analysis Criteria

The challenge and gap analysis applies several criteria for identifying cyber security research needs. The focus was either a lack of a specific resource (from the toolkit) or noted challenges faced in today’s fossil power plants. Extensive collaboration with the industry identified several challenges related to deploying and maintaining a cyber security program in the rapidly changing landscape. The below criteria were considered when identifying research gaps and future research needs:

- Lack of an industry recognized methodology, standard, or guidance.
- Absence of a technology to solve a cyber security challenge.
- A programmatic element that relies heavily on people to execute.

- Components that are not inherently secure. It is preferred that something is secure by design, and arrives in its most secured configuration, as opposed to requiring third party tools or manual steps to secure.
- High barrier of entry, either relating to capital cost, knowledge, and/or resource time to implement.
- Existing challenges in the industry.

Industry Challenges, Notable Gaps, and Research Needs

The research project held a developmental workshops and interviews with several utility members, researchers, and technology providers. The workshop and interviews identified many challenges from the industry, which helped shape the research framework.

Applying a Graded Approach to Cyber Security

Utilities need a graded approach to cyber security. The TAM and CCE hazard analysis approach allows an owner to identify and prioritize the most critical components. This means adding more layers to the defense in depth strategy for more critical, or susceptible components. The current “low-impact” regulatory requirements are a starting point for generating facilities. Many facility owners recognize that additional control measures are needed, particularly for critical components.

A graded approach also allows decision makers to perform a cost-benefit analysis and tradeoffs for purchases. Both vendors and technology developers, as well as end users, noted this as a challenge. Graded, risk-informed approaches offer many benefits. Notably, they allow for a more efficient, targeted approach to applying cyber security controls. The challenge is that a higher level of system understanding is needed to effectively deploy. It is also difficult to regulate, since audits (regulatory or otherwise) generally rely on a checklist of some fashion.

Research Area: Suggested areas of research to address this challenge include extending and refining the CCE, TAM, and Equipment Criticality methodologies through worked examples and industry pilots.

Asset and Configuration Management

Asset management has many definitions depending on the tool vendor or the user’s perspective. Asset management in the power plant refers to asset identification, spare parts, work management, and so on, using tools such as Maximo. Some facility asset management tools do not implement effective configuration management and change control. IT asset management refers to discovery and configuration management of networked IT equipment.

Asset management within the OT environment is a pivotal area of research as many other areas of cyber security rely heavily on effective asset and configuration management. The OT environment provides many challenges for typical IT asset management tools that would be used in a business environment, such as specific protocols for device families that may provide little asset information, islanded assets with no connection to the central OT environment, and reduced ability to perform active scanning. Because of these challenges, a business asset management system will generally perform poorly in the OT environment without considerable customization.

Configuration management refers to keeping plant installed equipment in configuration equilibrium. In configuration equilibrium, the documented digital asset configuration captured in facility documentation or “known good” baseline configuration agrees with the “as found” physical configuration. Configuration management is a key aspect of effective digital engineering, facility O&M, and cyber security. Asset and configuration management is the cornerstone of a cyber security program. It is needed to enable many cyber elements, such as patch management and backup and recovery planning.

IT asset management tools are widely used in other industries, as well as on the energy sectors business environment. OT asset management has many challenges that make it difficult to deploy the IT tools effectively. Some utilities do deploy tools, but they either take significant effort and expertise to deploy and maintain or provide limited value. The specific challenges include the following:

- Many assets are non-networked or part of a small, isolated network.
- Some OT assets are not discoverable by traditional IT asset management tools.
- OT equipment is sensitive to active network scanning.
- For tools that require installation of an agent, most OT equipment does not support installing an agent.
- Multiple different protocols are used that are not prevalent in IT.
- OT configuration changes are also not discoverable by traditional IT asset management tools.

Research Areas: Suggested areas of research to address these challenges include the following:

- Methods to record more useful information on device configurations, configuration and asset management for non-networked assets.
- Exploring the compatibility of OT protocols within traditional IT asset management tools.
- Developing formal testing guidance that are compatible with the limits and considerations of OT environments.
- Many networked devices within the OT environment will have vastly different configuration considerations, and this can lead to storing device configurations for separate device types in separate locations. Research needs to be conducted on the development of a central system for storing useful device configuration information, useful being specific information to that device family. With OT devices generally having nontraditional OS, there is little compatibility with currently available tools.
- Many OT devices are non-networked, meaning that there are few to no automated methods for gathering asset information for these devices. This provides the need for research-related to new methods to conduct asset discovery and management on non-networked devices, whether it is manual, automated, software-based, or hardware-based.

- Finally, OT networks are commonly viewed as being easily disrupted by cyber security tools. The final area of research to address this challenge is exploring commonly used OT protocols and active monitoring practices to extend the capabilities in automated asset management and developing best practices when extending the use of OT protocols and using active monitoring within an OT environment.

Changing Regulatory Landscape

North American fossil generation facilities must comply with the NERC CIP regulatory standards. Today, a majority of facilities fall into the NERC low-impact classification. There are also approximately 30–40% of the standards that are being revised or evaluated for a revision. Additionally, there are changing expectations with respect to how to comply with the standards. This was a noted challenge, and, in some cases, results in a focus on how to comply as opposed to how to deploy an effective cyber security strategy.

Research Area: Suggested areas of research to address this challenge include extending the CCE and TAM to develop compliance case studies to illustrate how the implementation of first principle methodologies lead to compliance.

Rapidly Changing Technologies and IT/OT Convergence

As discussed in Section 2, there are many life-cycle drivers. Technology developers and providers are striving to improve their capabilities and offerings to their customers to keep up with the cyber need. OT systems and components are increasingly built on IT infrastructure, allowing vendors to introduce IT cyber tools for use with their systems. Changing technologies like virtualization, cloud storage, and computing as a service are reshaping how OT vendors support their customers. In addition to the rapid pace of change, the knowledge and skill level needed to deploy and maintain these technologies are becoming a challenge.

Research Areas: Suggested areas of research to address this challenge include the following:

- Developing resource capabilities, knowledge, and skill sets needed for digital power plants of the future
- Partner with leading OT vendors to develop technology adoption curves that give facility owners insight into the technologies and skills needed in the future
- Developing resource knowledge and skill training guidelines that give facility owners guidance on available training options

Event Response

Event response, despite requiring swift action, is largely a manual effort. These manual efforts include evaluating alerts and alarms and developing and executing playbooks to respond and recover from an event. ICSs have the capability to create an overwhelming amount of information in the form of alerts and alarms, and it is the responsibility of the facility cyber team to filter through this information and identify abnormalities. While there are a number of automated IDSs that can help alleviate these issues (for those assets that are discoverable), the

cost and level of knowledge required to configure and maintain these systems can easily be cost or resource prohibitive. The facility cyber team is also tasked with the development of playbooks designed to respond to different types of events and, ultimately, execute the playbooks if the specific event arises.

Research Areas: Suggested areas of research to address this challenge include exploring new options in automated alarm filtering and automating actions as a result of alarming and integrating these capabilities with existing OT infrastructure. There is existing research in the area of artificial intelligence (AI) and confidence levels with regard to alarming. Additional research needs to be conducted to explore measuring confidence in identifying alarms that indicate cyber events. This confidence level would then be used to reduce the likelihood of false positive alerts. When an alarm is identified to indicate a cyber event, whether automatically or manually, there is a need to automatically act and respond according to an established playbook. This is an area worth researching, specifically due to the sensitive nature of OT equipment and its role within critical plant functions.

Another potential issue with cyber event response is unintended negative consequences. Traditional IT cyber event responses may cause a physical consequence such as automatic shutdown of a DCS server, when the DCS may still be needed to successfully place the plant in a stable and safe condition. Research that evaluates and provides guidance for safe cyber event responses as well as safe integration of IT cyber event response technologies can add significant value to the industry. A supporting research area is the development of tradeoff analysis that would identify consequence (CCE) for individual mitigations to inform the SOAR (playbooks, and so on) response. CCE could be used for network planning to enable designs that incorporate the tradeoff analysis and segment operations so that those with more consequence have less dependency.

Finally, there are existing efforts within the IT sector, and to prevent duplicate work, research could be conducted to explore applying IT practices related to automated alarm reporting and event response to OT infrastructure.

Secure Configurations for OT Assets

While some OT equipment is purpose-built, the majority include a number of features and functions that can be configured for a variety of purposes and functions. Because of this wide range of capabilities, OT equipment is typically built with an excess of features that will extend past the needs of a specific intended function. By default, a typical shipped configuration includes the majority of features enabled, mostly for ease of configuration. However, this is an issue because unused features that are enabled by default have the potential to provide attack pathways that can be exploited.

While there are vendors that provide guidance on device hardening, there is no standard deployed benchmarks for hardened configurations of OT equipment in the energy sector. This results in devices being deployed with unintended features and unused services enabled, allowing attack pathways, which would have been prevented with basic device hardening.

Research Areas: Suggested areas of research to address this challenge include the following:

- Developing an improved series of standards and libraries for device configurations, standardized documentation of security posture and countermeasures, guidance for vendor development and default hardening, and researching automated tools to support secure deployment.
- Extending the TAM CSDS concept to develop a library of Reference CSDSs for the energy sector to significantly improve efficiency across the entire sector. A library of Reference CSDSs would eliminate the redundancy of multiple utilities performing cyber assessments on the same equipment.
- There is currently little guidance for vendors related to the area of device hardening. Further research could include developing a series of best practices for vendors related to developing hardened configurations for equipment, making guidance for device hardening available, or shipping equipment with hardened configurations with only the necessary features enabled for the intended use.

Supply Chain Integrity and Testing

The supply chain represents a significant attack pathway; however, integrators and facility owners have little visibility into supply chain integrity vulnerabilities and limited options to secure the supply chain.

Integrators and facility owners evaluate vendors based on trust and compliance with programmatic requirements. Increasing the testing and evaluation capabilities within the supply chain processes would decrease the reliance on trust and programmatic requirements in the vendor assessment process.

Research Areas: Suggested areas of research to address this challenge include the following:

- Developing guidance and tools to increase testing capabilities within the supply chain. This includes new methods, standards, and best practices will need to be developed for vendors and facility owners. This research should include topics related to vulnerability discovery, vendor provided cyber information and considerations, and expectations for site implementation.
- Developing standard testing processes that include various testing methods used to discover vulnerabilities and presenting the findings in a standardized report.
- Developing vendor guidance on providing the purchaser relevant cyber program data and information. Examples of this information could consist of, but is not limited to, hardening configurations, Reference CSDSs, known good hash files, third-party software packages or libraries, and any included back-door accesses, whether enabled or not.

6

REFERENCES, BIBLIOGRAPHY, AND ABBREVIATIONS

References

1. *Changing Threat Landscape Study Report: Generation Cyber Security—Trends Across Cyber Security Incidents at Industrial Facilities*. EPRI, Palo Alto, CA: 2019. 3002015259.
2. Department of Homeland Security, Cybersecurity & Infrastructure Security Agency, “Ransomware Impacting Pipeline Operations,” Alert AA20-049A, <https://us-cert.cisa.gov/ncas/alerts/aa20-049a>.
3. The Process of Consequence-driven Cyber-informed Engineering (CCE), The Idaho National Laboratory (INL), 2018, INL/EXT-18-44814.
4. *Cyber Security Technical Assessment Methodology: Risk Informed Exploit Sequence Identification and Mitigation, Revision 1*. EPRI, Palo Alto, CA: 2018. 3002012752.
5. [*“SpaceX and Daring to Think Big.”*](#) Steve Jurvetson. January 28, 2015.
6. North American Reliability Corporation (NERC), “Lesson Learned: Loss of Automatic Generation Control During Routine Update” [20200403], April 14, 2020, https://www.nerc.com/pa/rrm/ea/Lessons%20Learned%20Document%20Library/LL20200403_Loss_of_AGC_During_Routine_Update.pdf.
7. *Risk Informed Target Level Topical Guide*. EPRI, Palo Alto, CA: 2019. 3002015760.
8. *Guideline for Establishing Criticality Rankings for Fossil Power Plant Systems and Components*. EPRI, Palo Alto, CA: 2012. 1025341.
9. IEC 61508 Ed. 2.0 b:2010 Functional Safety of Electrical/Electronic/Programmable Electronic Safety-Related Systems, Seven Parts.
10. IEC 61511 Ed. 1.0 b:2018, Functional Safety – Safety Instrumented Systems for the Process Industry Sector, Three Parts.
11. IEC 62443, Security for Industrial Automation and Control Systems, Four Parts.
12. Junda Zhu, Jae M. Yoon, David He, Yongzhi Qu, and Eric Bechhoefer, “Lubrication Oil Condition Monitoring and Remaining Useful Life Prediction with Particle Filtering,” *International Journal of Prognostics and Health Management* (2013): 1–15.
13. “Turbine Lube Oil System of 200MW,” presentation, <https://www.scribd.com/doc/84113887/Turbine-Lub-Oil-System>.
14. Sarah Freeman and Colleen Glenn, “Consequence Prioritization Impact Scenarios,” Cybercore Integration Center, *Idaho National Laboratory* (February 5, 2019): 31.

15. Wikipedia article on First Principles, In Physics, https://en.wikipedia.org/wiki/First_principle.
16. Aristotle, Physics 184a10–21.
17. Aristotle, Metaphysics 1013a14-15.
18. “[*Elon Musk's Mission to Mars*](#),” Chris Anderson, Wired.
19. *Cyber Security Technical Assessment Methodology, Risk Informed Exploit Sequence Identification and Mitigation, Revision 1*. EPRI, Palo Alto, CA: 2018. 3002012752.
20. NIST Special Publication 800-82, *Guide to Industrial Control Systems (ICS) Security*, Revision 2, May 2015.
21. Common Weakness Enumeration (CWE), <https://cwe.mitre.org/data/index.html>.
22. Common Vulnerabilities and Exposures (CVE), <http://cve.mitre.org/about/terminology.html>.
23. Common Vulnerability Scoring System (CVSS), <https://www.first.org/cvss/specification-document>.
24. National Vulnerability Database (NVD), <https://nvd.nist.gov/cwe.cfm>.
25. OWASP, Open Web Application Security Project website, https://www.owasp.org/index.php/Main_Page.
26. Pascal Meunier, *Classes of Vulnerabilities and Attacks*, Wiley Handbook of Science and Technology for Homeland Security, 2008.
27. Michael J. Assante and Robert M. Lee, *The Industrial Control System Cyber Kill Chain*, October 2015.
28. Stacey Cook, Sarah G. Freeman, and Curtis P. St. Michel, “CCE Phase 3: Consequence-based Targeting [INL-EXT-20-58090],” May 5, 2020.
29. *Cyber Security Data Flow Identification and Documentation Topical Guide*. EPRI, Palo Alto, CA: 2020. 3002015761.
30. IDART Information Assurance Red Team, <http://www.idart.sandia.gov>.
31. NIST SP 800-115, *Technical Guide to Information Security Testing and Assessment*, September 2008.
32. NIST Cybersecurity Framework Version 1.1, April 2018.
33. NIST *Framework for Improving Critical Infrastructure Cybersecurity*, Version 1.1, April 2018.
34. Sarah G. Freeman, Theodore Miller, and Curtis P. St. Michel, “*CCE Phase 4: Mitigations and Protections [INL-EXT-20-58091]*,” May 5, 2020.
35. Merriam-Webster, Online Dictionary, <https://www.merriam-webster.com/dictionary>.
36. STPA Handbook, Nancy Leveson, John Thomas, March 2018.
37. Security PHA Review for Consequence-Based Cyber Security, Edward Marszal, Jim McGlone, 2019.

38. North American Electric Reliability Corporation (NERC) Critical Infrastructure Protection (CIP) standards, <https://www.nerc.com/pa/Stand/Pages/CIPStandards.aspx>.

Bibliography

DHS Recommended Practice, Developing an Industrial Control Systems Cybersecurity Incident Response Capability

DHS Recommended Practice for Patch Management of Control Systems

EPRI, Access Control and Permission Management Guideline (3002014368)

EPRI, Control System Protocols and Scanning Guideline (3002014369)

EPRI, Cyber Security in the Supply Chain, Procurement Methodology Revision 2 (3002012753)

EPRI, Digital I&C Configuration Management and Hardening Guideline (3002011904)

EPRI, Incident Response Guideline (3002014147)

EPRI, Interactive Remote Access Guidance (3002011541)

EPRI, Patch Management Guideline (3002011187)

EPRI, Security Event Monitoring Guideline (3002014367)

NIST SP 800-63B, Digital Identity Guidelines, Authentication and Lifecycle Management

NIST SP 800-167, Guide to Application Whitelisting

Abbreviations and Acronyms

AD	active directory
AGC	automatic generation control
CCE	consequence-driven cyber-informed engineering
CCR	component criticality ranking
CIP	critical infrastructure protection
CML	control method library
CSDS	Cyber Security Data Sheets
CVE	common vulnerabilities and exposures
DCS	distributed control system
DOE	U.S. Department of Energy
DOS	denial of service
DTM	device type manager
EMS	energy management system
EPRI	Electric Power Research Institute

ER	Equipment Reliability Program
HART	Highway Addressable Remote Transducer
HCE	high consequence event
HMI	human-machine interface
I&C	instrumentation and control
ICS	industrial control system
IDART	Information Design Assurance Red Team
IDS	Intrusion Detection System
IEC	International Electrotechnical Commission
IT	information technology
INL	Idaho National Laboratory
LOPA	layer of protection analysis
MAC	media access control
NERC	North American Electric Reliability Corporation
NERC-CIP	NERC Critical Infrastructure Protection
NIST	National Institute of Standards and Technology
NVD	National Vulnerability Database
O&M	operations and maintenance
OCR	Operational Criticality Ranking
OS	operating system
OT	operational technology
PHA	process hazard analysis
PLC	programmable logic controller
RSDS	Relationship Set Data Sheet
RTF	run-to-failure
SCADA	supervisory control and data acquisition
SERP	System and Equipment Reliability Prioritization
SIEM	security information and event management
SIF	safety instrumented functions
SIL	safety integrity level
SIS	safety instrumented systems

SL	security level
SME	subject matter expert
SOAR	security orchestration automation and response
NIST-SP	NIST Special Publication
TAM	technical assessment methodology
TTP	tactics, techniques, and procedures

A

CYBER SECURITY TOOLKIT—EXAMPLES OF STANDARDS, GUIDES AND METHODOLOGIES, AND TOOLS AND TECHNOLOGIES

This cyber security toolkit is built for Life Cycle Risk Reduction Framework (LRRF) by providing examples of Standards, Guides, and Methodologies as well as sample list of available technologies. Its purpose is to categorize and group a sample of commonly used techniques and technologies that are utilized within the energy sector. The techniques and technologies are reviewed and then categorized as a methodology, standard, and guide, or a tool/technology.

This collection is not a comprehensive list of all tools that are currently available but instead a sample of tools available. There is an overload of available options and this can make it challenging to select appropriate tools. This cyber security toolkit aims at solving this challenge by providing a group of tools that are mapped to the LRRF to provide an idea of what type of tools can be used to fulfil specific needs.

Standards, Guides, and Methodologies

Technical Assessment Methodology

The TAM [4] is a risk informed cyber security engineering approach that provides a generation facility owner with confidence that the cyber risk has been reduced to an acceptable risk level that is also compliant with regulations. The TAM guides a user through a deliberate and repeatable process that efficiently converges attack surface identification and mitigation activities to an effective result. The TAM develops an installed configuration and data flow in which the attack surface is reduced as much as possible while still achieving functional objectives. Appropriate mitigating control measures are put in place to mitigate the attack surface vulnerabilities to an acceptable risk level. The TAM may be used at any point in the life cycles identified in Section 2.

Consequence-Driven Cyber-Informed Engineering

The CCE methodology is a process to identify HCEs that can be targeted for cyber attack by an adversary [3]. The CCE methodology is a four-phase process that includes, consequence prioritization, system-of-systems analysis, consequence-based targeting, and finally mitigations and protections. Consequence prioritization uses a scoring system to set a threshold of an acceptable HCE threshold and then prioritize all HCE's above the threshold. Next in the system-of-systems analysis, key information, access, and actions of an attacker are identified. In consequence-based targeting and system-of-systems analysis, the problem is approached from an adversary's point of view to identify where the control system is vulnerable. Finally, in the mitigation phase identified vulnerabilities are mitigated by control measures to an acceptable risk level.

System-Theoretic Process Analysis

System-theoretic process analysis (STPA) is a model that is used to identify accidents and hazards caused by unsafe and unsecure control actions within system control structures, including causal scenarios [36]. An evaluation is done on any causal scenarios that could lead to unacceptable losses and mitigations and controls are designed and implemented to eliminate the scenarios. This method of hazard analysis is gaining popularity in aircraft control system hazard analysis and other industries.

EPRI HAZCADS

EPRI is in the process of developing Hazards and Consequences Analysis for Digital Systems (HAZCADS) that integrates elements from System-Theoretic Process Analysis and Fault Tree Analysis (FTA) to identify digitally related components and system hazards. These hazards could manifest either through random failure, cyber security vulnerabilities, software failures, design flaws, and implementation errors. HAZCADS is applied systematically by; gathering plant design, system design, and hazard data, defining the losses, hazards, and the control structure, identifying any unsafe control actions, obtain or create a traditional fault tree, generate a System-Theoretic Fault Tree (SIFT), and finally allocate a target control method effectiveness score for each of the I&C system elements.

EPRI Component Criticality Ranking

The EPRI CCR guideline provides a summary of multiple methods for identifying and prioritizing critical components and systems within a power plant [8]. The methods described vary from simple to complex analysis that rank components relative to one another within the plant. The Low Complexity CCR Method relies on experienced plant personnel to review equipment within a master equipment list and rank the components based on their knowledge of the of the components purpose in a process, and the components functions. The Medium Complexity CCR Method is a more common method that asks the user a series of questions to determine the CCR. Typical questions can be grouped into critical or non-critical categories and relate to the overall operation of the plant or individual systems. Finally, the High Complexity CCR Method is the most involved method and it involves an algorithm-based approach that takes several inputs to formulate a component CCR.

Cyber Process Hazards Analysis

Security PHA Review for Consequence-Based Cyber Security (SRP) is a systematic consequence-driven approach that utilizes several standards related ICS to conduct cyber security risk assessments within the ICS environment [37]. A Cyber PHA is conducted in a series of steps that commonly include an assessment, an analysis, a workshop, and finally crafting a mitigation plan. The plant systems are analyzed to identify existing vulnerabilities, which are then evaluated to determine whether any of the vulnerabilities could be exploited during a cyber event. A Cyber PHA workshop is then conducted and a series of threat scenarios are crafted using the discovered vulnerabilities. Upon completion of the workshop a comprehensive report for the existing vulnerabilities and mitigation strategies is created.

North American Electric Reliability Corporation CIP

The NERC is responsible for maintaining the critical infrastructure protection (CIP) standards for the bulk electric system (BES) in North America [38]. NERC CIP compliance is required US energy utilities. The standards focus on improving the North American power system security. Currently there are eleven standards that are currently enforceable with several new and updated standards pending future enforcement. First, BES's are labeled as high, medium, or low impact which is outlined in NERC CIP-002. The remaining standards (NERC CIP 003-014) list requirements for high, medium, and low impact BES entities.

IEC 62443

The IEC is in the process of releasing an IEC 62443 [11] series of standards for ICS cyber security. The series covers a range of scenarios from component to system security, vendor development life cycle security, to security associated with technical services. IEC 62443 is also developing a standard for cyber security risk levels.

ISO 27001

The International Standards Organization (ISO) maintains ISO 27001, a leading standard in information security management systems. IS27001 is intended to be a systematic cost-effective approach to secure an organizations information. The chief focuses of ISO 27001 are ensuring only authorized users have access, only the authorized user has access to the information, and information must be available to authorized users when needed. Organizations can use the standard to aid in the implementation of processes to meet regulatory requirements, or to receive an ISO 27001 certification through an audit process.

NIST Cyber Security Framework

The NIST maintains the Cybersecurity Framework (CSF), a voluntary series of standards developed for critical infrastructure within the United States. NIST CSF is a cost-effective approach for creating prioritized, flexible, performance based, and repeatable measures and controls that aid critical infrastructure in identifying, assessing, and managing cyber security risks. This framework may be used in the initial designing of an information security system, assessing the current state of an information security system, or updating specific processes.

NIST SP 800

The NIST Special Publication (SP) 800 series guidelines are developed with the intention to aid in the compliance with FISMA Federal Information Processing Standards (FIPS). Relevant NIST SP 800 standards include 800-53, 800-82, 800-83. In addition to the aforementioned standards there are standards for information security practices related to many technologies that are emerging in the energy sector, thus they are a good source for initial research in standard practices for new or soon-to-be adopted technologies.

IEC 61508 and IEC 61511

The IEC 61508 Functional Safety of Electrical/Electronic/Programmable Electronic Safety-related Systems [9] and IEC 61511 Ed. 1.0 b:2018, Functional Safety – Safety Instrumented Systems for the Process Industry Sector, Three Parts [10], sets the guidelines for SIS and safety system Safety Integrity Levels (SIL). These guidelines include direction on the systems design, implementation, and operation. SIL is broken down to four levels, with SIL4 being the highest risk. The levels are assigned to systems based on the risk involved with operating that specific system. The standard is comprised of seven parts including (1) general requirements, (2) requirements for electrical/electronic/programmable electronic safety-related systems, (3) software requirements, (4) definitions and abbreviations, (5) examples of methods for the determination of safety integrity levels, (6) guidelines on the application of IEC 61508-2 and IEC 61508-3, and (7) Overview of techniques and measures.

CISA Assessments and CSET

The Cybersecurity and Infrastructure Security Agency (CISA) is an agency within the Department of Homeland Security, and they offer to conduct free security assessments for both federal and private organizations. These free assessments include a self-assessment tool (Cybersecurity Evaluation Tool <https://www.us-cert.gov/ics/Downloading-and-Installing-CSET>) that is downloaded and guides the user through an assessment, onsite field assessments, network design architecture reviews, and network traffic analysis and verification. The information gained from the assessment is then provided back to the organization with additional context that aids in understanding and developing effective defense-in-depth processes.

Cybersecurity Capability Maturity Model

Cybersecurity Capability Maturity Model (C2M2) is a free program released by the DOE to perform a self-evaluation of your organizations cyber security maturity. The tool walks the user through perforating a self-evaluation by asking specific questions from a full range of topics related to cyber security readiness. Once the questions are completed a report is generated that displays gaps and deficiencies in the organization's current cyber security readiness. At this point the organization may prioritize and plan to address the identified deficiencies.

Technology Capabilities

This section provides a summary of a limited list of technology capabilities that are available to fossil generation plant owners. Note that this list is presented as an example of the multitude of technology capabilities and is not a comprehensive list of technology capabilities and vendors. In fact, the rapidly changing security environment, technology advancements, and new market entrants make it impossible to present a comprehensive list here.

The proliferation of tools and rapid pace of functionality in tools to manage systems and automate processes is mind boggling. And power generation owners have limited resources to research trends, become educated on technologies, and evaluate & select products.

There are industry analyst companies that provide research and consulting on market and vendor trends such as Gartner and Forrester. These are valuable resources when evaluating tool and automation strategy and can be utilized through subscription or individual purchases. Your company most likely has Gartner, Forrester, or other analyst subscription, and if so, make full use of the resource. If not, work with your executive sponsor to obtain the research. Vendors will often pay to publish reprints of industry analyst reports that provide favorable analysis and will make them available to potential customers.

Gartner Research provides consulting and research reports through Gartner Magic Quadrant and Critical Capabilities Reports (<https://www.gartner.com/en/research/methodologies/magic-quadrants-research>).

Forrester also provides consulting and publishes market trends and vendor research reports. The Forrester Wave and other research reports provide insight into vendor capabilities for specific markets (<https://go.forrester.com/research/>).

Hardware Integrity Checking

Hardware integrity checking consists of a series of checking mechanisms to ensure that newly acquired hardware has not been tampered with. Typically, a chip that detects tampering or an algorithm that runs based on the physical characteristics of the hardware to produce a “signature” that is able to detect tampering. Tamper seals also provide tamper detection. This is new area of research that has not yet yielded enterprise level products for everyday use. Physically Unclonable Functions (PUFs) and Trusted Platform Modules (TPM) are in use or in development. This area is critical to supply chain security and we expect to see significant new research into this area.

Penetration Testing

Penetration testing is the process of assessing security practices and policies by circumventing security features based on knowledge of the networks and utilizing tools commonly used by attackers. These tests can be a part of a known overt test, or part of a covert test with as little as upper management knowledge. Not only will penetration testing evaluate the network security and vulnerabilities, but it will also test response policies in place. Penetration testing can be done via automated testing programs or performed live by security professionals either within your organization or by external vendors.

Examples include Netsparker Security Scanner, Wireshark, and Metasploit.

Intrusion, Prevention, and Anomaly Detection Tools IDS/IDPS/HIDS

Intrusion Detection and Prevention Systems (IDPS) utilizes network taps or span ports to evaluate a copy of the network traffic at their location. Network tap or span port implement passive scanning with little impact on network traffic. IDSs can detect network anomalies with a variety of different methods ranging from hard coded rules based on network ports, to advanced AI models trained on historical network traffic. IDPSs also have the ability automatically respond to certain anomalies to “prevent” the intrusion. Many DCS providers also include IDS and Asset Management platforms in their offering.

Examples include:

- Security Events Manager from Solar Winds, Falcon from CloudStrike, and Snort
- Commercial IDS tools from Claroty, CyberX, Dragos, Forescout, and Nozomi
- Networking equipment with IDS Capability from Cisco, Fortinet, and Palo Alto
- Many Host Based Solutions

Security Information and Event Management

A Security Information and Event Management (SIEM) systems aggregate data from multiple sources so that the data may be analyzed and used to catch abnormal behaviors within the network. Data within a SIEM is normalized, so that regardless of sources all data should be clearly relatable. Rules are then developed and tuned to provide meaningful alarming and reporting. After an alarm, logging from data aggregation can then be used for forensics.

Examples include AT&T AlienVault, IBM QRadar, Gavwell, Graylog, RSA NetWitness, LogRhythm, McAfee SIEM, and Splunk.

Event Response and Automation

Security Orchestration Automation and Response (SOAR) is a robust extension of SIEM technology and is the process of performing automated tasks within several programs or processes to complete a complex security action. This often includes aiding responders in decision making, such as providing response playbooks, and managing the event. Orchestration technologies allow for complex security actions to be performed with little or no user input. These automated tasks are typically performed on demand, scheduled at fixed intervals, or driven by events. Another important consideration is that complex tasks that span over several programs are prone to misinterpretation, skipped steps, or provided incorrect inputs when performed by a user. Available technologies are typically limited to specific software suites, OS, or complexity of interactions.

Examples include Rapid7, and Helix from FireEye.

Identity and Access Management

Identity and Access Management (IdAM) is comprised of policies and technologies that ensure users have appropriate access to systems. This is achieved by authenticating users and verifying the identity of users and hardware endpoints. Authenticating user credentials and ensuring the user has sufficient privilege to access the systems are handled by several different technologies with AD being the most common for MS Windows networks.

Examples include AD with MS Windows Server, RSA Identity and Governance, SailPoint IdentityIQ.

Secure Remote Access

Secure remote access is a secure, encrypted connection to a network from another geographic or logical location. A remote connection to the corporate network may be made from company equipment, or from user equipment if permitted by company policy. Regardless of company or employee owned, the equipment must be maintained and properly secure with company standards. Segmenting networks into DMZs and using VPN are important aspects of the architecture.

Examples include:

- Secure Remote Access solutions from EPRI or Claroty
- Hardware or Software VPNs from MicroSoft, Cisco, Fortinet, or Palo Alto

Network Boundary Protection and Network Access Control (NAC)

A site's network is segmented into several sub-networks. All external connections to these networks, which include connections between networks as well as external access, need boundary protection. Boundary protection ensures that only authorized machines, protocols, messages, and users are allowed access as well as limiting access to only authorized segments or machines.

Examples include:

- Firewalls from Cisco, Palo Alto or Jupiter
- Network Switches from Cisco, Allen Bradley, or Moxa
- Routers from Cisco, Moxa, or Hirschman
- NAC software solutions from Impulse SafeConnect or SolarWinds

Asset and Configuration Management

Asset and configuration management is the practice of storing and managing information related to site assets. This information includes but is not limited to, asset identifying information, physical location, network addresses, firmware versions, installed software, service information, configuration files, backups, and relationships to other systems. This information can be acquired and maintained from existing asset or equipment databases, network traffic capture, or local agents. Management of Change (MOC) is a critical component that requires a known good baseline configuration and rigorous process to manage changes. Many of the above IDPS and SIEM solutions include capabilities to identify assets and changes.

There is a significant distinction between power plant asset management and IT asset management that can create confusion.

Examples include:

- Power plant Asset Management Systems from Maximo, SAP, or Passport
- Facility Asset Management Systems from Fiix, emaint, or Asset Essentials
- IT Asset Management Systems from ABB, Symantec, ManageEngine or Asset Panda
- IT/OT Asset Discovery and Management tools from Tenable/Indegy, Langner OT Base, or CyberX

Computerized Maintenance and Materials Management

Computerized Maintenance and Materials Management (CMMS) tools are utilized to schedule and manage maintenance activities including spare parts and warehouse activities to increase asset life and reduce downtime. CMMS tools are often combined with facility Asset Management. CMMS tools rely on information from asset management systems, and uses the information related to assets maintenance requirements, affected systems, and criticality to develop optimized maintenance strategies. These strategies are ultimately developed based on the goals of the system, with goals typically related to uptime, reliability, and life cycle management.

Examples include:

- Power plant CMMS from Maximo, SAP, or Passport
- Facility CMMS tools from Fiix, emaint, or Asset Essentials

Compliance Management

Corporate Compliance and Oversight (CCO) are used to centralize regulatory compliance management and reporting. Typical technologies for compliance management are purpose-built applications for specific reporting needs, or local processes and reporting documents. These practices alleviate the typical issues revolving around compliance management including, duplicate efforts, and inconsistencies with recording. Some tools are purpose built for specific regulations such as HIPPA or NERC CIP, while others can be configured to manage multiple regulatory frameworks in a single tool.

Examples include LogicGate Risk Cloud, Spirion, and NetWrix.

Encryption and Data Protection

Encrypting data at rest and in transit is an effective control measure to protect against theft or alteration of data. Sensitive information such as usernames and passwords, device configurations and backups should be encrypted when stored or transmitted over a network. With cloud storage providers becoming a standard practice, it is important to ensure that transmittal and storage employ encryption best practices meet your organizations data protection requirements. Encryption is typically a technology embedded in a larger tool or software that uses an algorithm such as RSA or AES 256 to perform the encryption calculation.

Examples include:

- Data in-transit algorithms such as TLS, HTTPS, or SSH
- Data at-rest tools such as Bitlocker, Protegrity, IBM Guardium

Software and File Verification

Software verification can be accomplished by several methods. One common approach is creating a checksum or HASH for all guarded files for a “known good” baseline is an essential method for verifying software and file integrity. Software downloaded from a website or vendor can be verified by running the HASH algorithm and comparing the result to the known good result from the vendor. Files can also be scanned and verified using the same method and can be automated. Scans can be done in response to events, periodically on a schedule, or performed

real-time. The frequency will depend on the security risks and needs as well as resources available. Checksums and HASHs are useful even in the event when scans will never be performed as described earlier, as they can be used to determine changes that may have been made unintentionally or identify files that were damaged from any means malicious or not.

Another common method for verifying software or files is via a code signing certificate or software publisher certificate. A cryptographic public-private key pair is generated. The entity developing (and/or publishing) the code, holds the private key, while a third party (typically a certificate authority or CA) holds the public key. The CA verifies the entity is bound to that public key. The software developer then signs the code with the private key, which can then be verified by the user through the CA and associated public key.

Examples include hashing algorithms such as SHA-1, SHA-2, MD5, HMAC-SHA.

Backup and Recovery

In the case of an event, having backups for system and component baselines is essential for recovery plans. In addition to creating, storing, and protecting backups, there need to be processes for recovery from the backups. Various technologies can perform backups when requested or scheduled on set frequencies. Most automated backup technologies are associated with MS Windows or Linux networks or transactional databases. Automated backup technologies will capture a disk image, specific directories and files, or transactions. When selecting a backup technology, ensure that the recovery processes satisfy your organizations security requirements with regards to data recovery, and recovery processes.

Examples include enterprise network solutions such as Commvault, Data Protection Suite from Dell, and Spectrum Protect from IBM.



Export Control Restrictions

Access to and use of this EPRI product is granted with the specific understanding and requirement that responsibility for ensuring full compliance with all applicable U.S. and foreign export laws and regulations is being undertaken by you and your company. This includes an obligation to ensure that any individual receiving access hereunder who is not a U.S. citizen or U.S. permanent resident is permitted access under applicable U.S. and foreign export laws and regulations.

In the event you are uncertain whether you or your company may lawfully obtain access to this EPRI product, you acknowledge that it is your obligation to consult with your company's legal counsel to determine whether this access is lawful. Although EPRI may make available on a case by case basis an informal assessment of the applicable U.S. export classification for specific EPRI products, you and your company acknowledge that this assessment is solely for informational purposes and not for reliance purposes.

Your obligations regarding U.S. export control requirements apply during and after you and your company's engagement with EPRI. To be clear, the obligations continue after your retirement or other departure from your company, and include any knowledge retained after gaining access to EPRI products.

You and your company understand and acknowledge your obligations to make a prompt report to EPRI and the appropriate authorities regarding any access to or use of this EPRI product hereunder that may be in violation of applicable U.S. or foreign export laws or regulations.

The Electric Power Research Institute, Inc. (EPRI, www.epri.com) conducts research and development relating to the generation, delivery and use of electricity for the benefit of the public. An independent, nonprofit organization, EPRI brings together its scientists and engineers as well as experts from academia and industry to help address challenges in electricity, including reliability, efficiency, affordability, health, safety and the environment. EPRI also provides technology, policy and economic analyses to drive long-range research and development planning, and supports research in emerging technologies. EPRI members represent 90% of the electricity generated and delivered in the United States with international participation extending to nearly 40 countries. EPRI's principal offices and laboratories are located in Palo Alto, Calif.; Charlotte, N.C.; Knoxville, Tenn.; Dallas, Texas; Lenox, Mass.; and Washington, D.C.

Together...Shaping the Future of Electricity

Program:

Cyber Security for Generation Assets

© 2020 Electric Power Research Institute (EPRI), Inc. All rights reserved. Electric Power Research Institute, EPRI, and TOGETHER...SHAPING THE FUTURE OF ELECTRICITY are registered service marks of the Electric Power Research Institute, Inc.

3002019700

Electric Power Research Institute

3420 Hillview Avenue, Palo Alto, California 94304-1338 • PO Box 10412, Palo Alto, California 94303-0813 USA
800.313.3774 • 650.855.2121 • askepri@epri.com • www.epri.com