Operational Technology Preparedness

A Risk-Based Safety Approach to Scoping Security Tests for Cyber Incident Response and Recovery

Alexander James Staves

School of Computing and Communications
Lancaster University

This dissertation is submitted for the degree of
Doctor of Philosophy

Graduate College August 2023

To my mother, for your unconditional love and support. Thank you for believing in me.

"We reject techniques like torture regardless of whether they're effective or ineffective because they are barbaric and harmful on a broad scale. It's the same thing with cyber warfare. We should never be attacking hospitals. We should never be taking down power plants unless that is absolutely necessary to ensure our continued existence as a free people."

- Edward Snowden

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the list of Publications and Acknowledgements.

Alexander James Staves
August 2023

Acknowledgements

I would first like to acknowledge my supervisors, Dr. Benjamin Green, Dr. Antonios Gouglidis, and Prof. David Hutchison, for their continuous support and guidance throughout the years. Ben, without your passion for ICS/OT cyber security, my research would not have been possible at Lancaster University. Your work in developing the ICS testbed and our many impromptu meetings together made my time here significantly more enjoyable. Antonios, not only did I have the pleasure of having you as a supervisor for my PhD, but I also owe you many thanks for introducing me to the world of ICS/OT as my supervisor for my undergraduate dissertation. Finally, David, while we may not have been able to meet in person as much as I had hoped, I greatly appreciate your patience and many words of wisdom during my time here. I hope I have resolved any misspellings of your name within the thesis.

To Frazer-Nash, thank you for giving me the opportunity to take part in research that will be applied in industry and for shaping the direction of my PhD. This opportunity and the funding provided by it have contributed significantly to the development of my knowledge and skills. I also owe a lot of gratitude to the numerous participants in my studies, which led to several engaging and thought-provoking conversations.

To my parents, words cannot describe how much I appreciate your support over the years, from travelling thousands of miles to visit me for a few days to reading my thesis from start to finish. I would also like to thank my brothers, who contributed significantly to a healthy work-life balance with our countless concert outings and late-night calls.

To my "other" family, the doctors-to-be of B55 and B59, our daily office debates ranging from philosophy to mathematics have kept me thoroughly entertained throughout the years. While my cholesterol levels may not thank you for our trips together to the various fast food outlets on campus, I do.

Abstract

Following the advent of Industry 4.0, there have been significant benefits to industrial process optimisation through increased interconnectivity and the integration of Information Technology (IT) and Operational Technology (OT). However, this has also led to an increased attack surface for cyber threat actors to target. A growing number of cyber attacks on industrial environments, including Critical National Infrastructure, has, subsequently, been observed. In response, government and standardisation organisations alike have invested considerable resources in improving the cyber security of these environments. This includes response and recovery, often used as a last line of defence against cyber attacks. However, due to the unique design philosophies of Industrial Control Systems (ICS), several challenges exist for effectively securing these systems against digital threats.

Through an analysis of standards and guidelines, used for assessing and improving cyber incident response and recovery capabilities, and stakeholder engagement on the implementation of these in practice, this thesis first identifies the challenges that exist when it comes to preparing for cyber incidents targeting ICS/OT environments. In particular, risk management, which involves identifying, evaluating, and prioritising risks and finding solutions to minimise, monitor, and control these, was found to be essential for improving preparation for cyber incidents. Assurance techniques are used as part of risk management to generate evidence for making claims of assurances about security. Alongside this, adversary-centric security tests such as penetration tests are used to evaluate and improve cyber resilience and incident response capabilities by emulating the actions of malicious actors. However, despite the benefits that these provide, they are currently not implemented to their full potential due to the safety and operational risks that exist in ICS/OT environments.

This thesis contributes to academic and industry knowledge by proposing a framework that incorporates methods for identifying and quantifying the safety and operational risks of conducting adversary-centric security tests within ICS/OT environments. In understanding the risks, these engagements can be scoped using precise constraints so as to maximise the depth of testing while minimising risk to safety and the operational process. The framework is then evaluated through a qualitative study involving industry experts, confirming the framework's validity for implementation in practice.

Publications

The following publications were used, in part verbatim, in the writing of this thesis:

Staves, A., Balderstone, H., Green, B., Gouglidis, A., and Hutchison, D. (2020). *A framework to Support ICS Cyber Incident Response and Recovery*. In the 17th International Conference on Information Systems for Crisis Response and Management, pages 638-651.

Miller, T., **Staves, A.**, Maesschalck, S., Sturdee, M., and Green, B. (2021). *Looking Back to Look Forward: Lessons Learnt from Cyber-Attacks on Industrial Control Systems*. International Journal of Critical Infrastructure Protection, Volume 35, Article 100464.

Staves, A., Anderson, T., Balderstone, H., Green, B., Gouglidis, A., and Hutchison, D. (2022). *A Cyber Incident Response and Recovery Framework to Support Operators of Industrial Control Systems*. International Journal of Critical Infrastructure Protection, Volume 37, Article 100505.

Staves, A., Gouglidis, A., and Hutchison, D. (2022). *An Analysis of Adversary-Centric Security Testing within Information and Operational Technology Environments*. Digital Threats: Research and Practice.

Maesschalck, S., **Staves, A.**, Derbyshire, R., Green, B., and Hutchison, D. (2023). *Walking under the Ladder Logic: PLC-VBS: a PLC Control Logic Vulnerability Scanning Tool*. Computers & Security, Volume 127, Article 103116

Staves, A., Gouglidis, A., and Hutchison, D. (2023). *Risk-Based Safety Scoping of Adversary-Centric Security Testing on Operational Technology*. (Under Review)

Table of contents

List of figures

List of tables

Nomenclature

Acronyms / Abbreviations

AI Artificial Intelligence

ATT&CK Adversarial Tactics, Techniques, and Common Knowledge

CAF Cyber Assessment Framework

CIA Triad Confidentialy, Integrity, Availability Triad

CISA Cybersecurity and Infrastructure Security Agency

CIS Center for Internet Security

CKC Cyber Kill Chain

CNI Critical National Infrastructure

CREST The Council for Registered Ethical Security Testers

CSC Critical Security Controls

CSIRT Cyber Security Incident Response Team

CVE Common Vulnerability and Exposure

DiD Defence in Depth

DMZ Demilitarized Zone

DNS Domain Name System

DoS Denial of Service

DWI Drinking Water Inspectorate

ENISA European Union Agency for Cybersecurity

EU European Union

FTA Fault Tree Analysis

GDPR General Data Protection Regulation

HAZOP Hazard and Operability

HMG His Majesty's Government

HMI Human Machine Interface

HSE Health and Safety Executive

HVAC Heating, Ventilation, and Air Conditioning

I/O Input/Output

IAEA International Atomic Energy Agency

ICS Industrial Control System

IEC International Electrotechnical Commission

IoT Internet of Things

IP Internet Protocol

ISO International Organization for Standardization

IT Information Technology

MCS Minimal Cut Set

NCSC National Cyber Security Centre

NDA Nuclear Decomissioning Authority

NEI Nuclear Energy Institute

NERC North American Electric Reliability Corporation

NHS National Health Service

NIS-D Network and Infromation Systems Directive

NIST National Institute of Standards and Technology

NPSA National Protective Security Authority

NRC Nuclear Regulatory Commission

NSS Nuclear Security Series

NVD National Vulnerability Database

OG Operational Guidance

ONR Office for Nuclear Regulation

OS Operating System

OT Operational Technology

P&I Piping and Instrumentation

PERA Purdue Enterprise Reference Architecture

PLC Programmable Logic Controller

R&R Response and Recovery

RCE Remote Code Execution

RQ Research Question

RTT Round Trip Time

RTU Remote Terminal Unit

SCADA Supervisory Control And Data Acquisition

SIS Safety Instrumented System

SOC Security Operations Centre

SP Special Publication

SSL Secure Sockets Layer

STIX Structured Threat Information eXpression

SyAPs Security Assessment Principles

TAG Technical Assessment Guide

TCP Transmission Control Protocol

TiDICS Testing in Depth for ICS

TLS Transport Layer Security

TTP Tools, Techniques and Procedures

UK United Kingdom

USA United States of America

Chapter 1

Introduction

During the past decade, many technological advances have been observed for the continuous improvement of the industrial process, otherwise known as Industry 4.0. A continued drive for interconnection within these environments has been prioritised due to the benefits it yields to areas including process optimisation, energy efficiency, proactive maintenance and more. The introduction of widely adopted protocols to support interconnectivity (i.e. TCP/IP) within this context can be considered the most significant technical evolution of recent years, while also the greatest catalyst of risk; resulting in a significant growth of attack surface for adversaries to target. While cyber attacks are not a novel concept in themselves, an increasing amount of these targeting industrial environments specifically, including Critical National Infrastructure (CNI), have been observed recently. Despite recent significant improvements for securing industrial assets against digital threats, cyber incident response and recovery still remains an essential part in all cyber security strategies; providing a last line of defence to minimise impact to operations in the event that security measures fail.

This chapter introduces the concept of the underlying technology within industrial environments, known as Industrial Control Systems and Operational Technology (ICS/OT), its application within CNI, Industry 4.0 and its effect on ICS/OT cyber security, and the challenges of effectively responding to and recovering from cyber attacks; forming core research motivation and objectives for the thesis. The focus of the research throughout this thesis is on preparation for responding to and recovering from cyber attacks on ICS/OT; more specifically, the challenges that Industry 4.0 presents for this and areas for improving it.

1.1 Research Motivation

1.1.1 Operational Technology and Critical National Infrastructure

Operational Technology is a term that has gained popularity over the last fifteen years to describe "hardware and software that detects or causes a change, through the direct monitoring and/or control of industrial equipment, assets, processes and events", as defined by Gartner [84]. They are designed with the intent to view, monitor and/or control a physical process, as opposed to traditional Information Technology (IT) which is designed to store, transfer and manipulate information (i.e. data). Some examples of OT include:

- Factory automation systems: used to control and monitor production lines.
- Supervisory Control and Data Acquisition (SCADA) systems: used to control and monitor industrial processes.
- Building management systems: used to control and monitor heating, ventilation, and air conditioning systems in buildings.
- Traffic management systems: used to control and monitor traffic flow in cities.
- Power grid control systems: used to control and monitor the distribution of electricity.

Different types of OT can also be categorised based on their function within an industrial process. The Purdue Model [59], an extended version of this being illustrated in Figure 1.1, is a reference architecture used for the hierarchical separation of systems within an industrial context based on their design intent. Within each level of this model can be found different types of OT that serve specific functions within the overall operational process.

Within the safety zone lies systems, known as Safety Instrumented Systems (SIS), which are critical for the safe operation of industrial environments. These systems are designed to prevent accidents and protect against hazards in industries such as chemical, oil and gas, and nuclear power [81] and are used in situations where the consequences of a failure or malfunction can be severe, such as the release of toxic chemicals or a nuclear accident. SIS are typically composed of sensors, controllers, and final control elements, such as valves and switches, which are used to monitor a process and take automatic corrective action if a hazardous situation is detected. These systems are used to return environments to a safe state in the event that a loss of safety occurs. For example, if the temperature or pressure in a chemical plant exceeds safe limits, a SIS will automatically shut the process down to protect workers and assets.

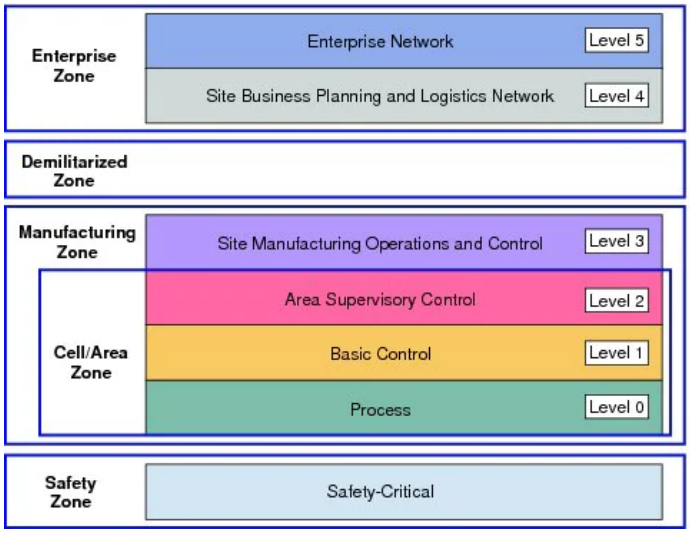

Fig. 1.1 The Extended Purdue Enterprise Reference Architecture [59]

Within the process level (level 0) lies devices that are commonly referred to as sensors and actuators [160]. Sensors are devices that detect changes in physical conditions, such as temperature, pressure, or motion, and convert them into a usable form, such as an electrical signal. Actuators, on the other hand, are devices that receive input signals and use them to control physical processes, such as opening and closing valves or moving mechanical parts.

The basic control level (level 1) is where localised control functions are performed. This level is typically composed of Programmable Logic Controllers (PLCs) or other ICSs that use the input from the process level to make decisions and control processes [22]. For example, in a factory, the control level may include PLCs that are used to monitor the output of sensors on the production line and make decisions about how to control the process, such as adjusting the speed of a conveyor belt or turning on and off specific actuators.

The supervisory and site manufacturing operations and control levels (level 2 and 3) are responsible for overseeing and coordinating the activities of the cell/area zone and manufacturing zone respectively. These levels are typically composed of SCADA systems or other systems such as engineering workstations or Human Machine Interfaces (HMI) that are used to monitor and control industrial processes [24]. For example, the supervisory level may include a SCADA system that is used to monitor the output of several basic control areas and make decisions about how to coordinate the activities of the different control systems in the factory.

In the context of Critical National Infrastructure (CNI), OT systems are essential for ensuring the reliability and safety of essential services, such as the electricity grid, water and wastewater systems, and transportation networks. These systems are often used to control

critical infrastructure, such as power plants and factories, and a breach in their security can have serious consequences [147]. For example, a cyber attack on a power grid control system could cause widespread power outages, disrupting essential services and potentially causing physical damage. CNI is defined by the National Protective Security Authority (NPSA) as the "facilities, systems, sites, information, people, networks and processes that are necessary for a country to function and upon which daily life depends" [205]. The UK government considers thirteen sectors as CNI and these are chemicals, civil nuclear, communications, defence, emergency services, energy, finance, food, government, health, space, transport, and water. CNI is typically managed by government agencies or regulated industries, and have recently become more and more targeted by cyber threat actors [171]. As such, ensuring the security of CNI has become a high priority for many governments. This involves implementing robust security measures, such as regular security assessments and incident response plans, to protect against cyber threats and other potential disasters [187, 196].

1.1.2 Industry 4.0 and ICS/OT Cyber Security

Industry 4.0, also known as the fourth industrial revolution, is a term that refers to the integration of advanced technologies, such as the Internet of Things (IoT), artificial intelligence (AI), and automation, into the manufacturing and industrial sectors [239]. The goal of Industry 4.0 is to create a more connected, efficient, and flexible manufacturing environment that can respond to changing market demands and deliver higher-quality products. This includes the use of process optimisation through means such as remote telemetry and machine learning algorithms to analyse data and make timely decisions for making changes to the production process [220]. In addition, industry 4.0 also involves the integration of OT and IT systems, which allows for greater collaboration and coordination between different parts of an environment [215].

While Industry 4.0 has brought about several technological advancements for increasing the efficiency of industrial environments, numerous challenges relating to this have also emerged. As industrial organisations adopt more connected and intelligent systems, the potential attack surface for cyber threat actors grows more significant and complex [147]. At the same time, these threat actors are constantly evolving their techniques and tactics, making it difficult for organisations to keep up and protect against new threats [171]. Organisations have, therefore, developed the need to adopt a proactive and adaptive approach to ICS/OT cyber security, when historically, cyber security was not prioritised [34]. This involves implementing advanced technologies and tools, such as AI and machine learning, to help detect and respond to potential threats in real-time [141]. In addition, investing in continuous

monitoring and assessment of ICS/OT to identify and address potential vulnerabilities before cyber criminals can exploit them is essential [148].

The reliance on novel software and technology used within these environments also presents several challenges. As more and more processes are automated, there is a greater need for robust and secure software to control these processes. This can be difficult to achieve, as ICS often have unique requirements and constraints that make it challenging to develop and implement secure software [75].

Overall, the impact of Industry 4.0 on ICS/OT cyber security has not gone unnoticed in recent years [72, 213]. As industrial organisations adopt more connected and intelligent systems, they must also invest in robust ICS/OT cyber security measures and practices to protect against the potential risks and vulnerabilities that these technologies can introduce. Implementing defensive measures through proactive and adaptive approaches, have become essential for organisations to protect their operations and critical infrastructure from the growing threat of cyber attacks targeting ICS/OT.

1.1.3 Cyber Incident Response and Recovery for ICS/OT

One of the most crucial aspects of any organisation's cyber security strategy is cyber incident response and recovery (R&R) [213]. Acting as a last line of defence in the event that a cyber attack succeeds, the goal of R&R is to contain the threat, prevent further impact on business continuity, restore affected systems to their normal state of operation and reduce the impact caused [187]. While effective R&R for traditional IT environments is important to prevent a loss of capital and reputation, R&R for ICS/OT environments is essential as cyber attacks targeting these can lead to a loss of safety, including the damage of expensive equipment or resources and, more importantly, human injury or death. For example, in 2015 and 2016, several Ukrainian electricity distribution companies were targeted by specialised malware, leading to power outages across major cities, including Kyiv, for several hours [287, 62]. While the impact of these attacks was minimised, had response and recovery actions failed, this could have led to essential services such as communication and health services becoming incapacitated. Another notable attack, caused by the Stuxnet malware in 2010, exploited Siemens software using four zero-day vulnerabilities [199]. Through this, the centrifuge speeds of the several uranium enrichment facilities in Iran were modified, causing an operational shutdown, destroying several thousand machines, and leading to considerable delays in the development of the Iranian Nuclear Programme [74].

To ensure consistent control of CNI security, governmental bodies have started implementing several strategies. Notably, the European Union (EU) approved the introduction of the Network and Information Systems (NIS) Directive in 2016 into national laws of all

of its member states [72]. This was the first piece of EU-wide legislation on cyber security that was introduced with the aim of achieving effective cyber security capabilities across all member states. While all member states adopted the NIS Directive, its implementation was deemed too complicated and resulted in inconsistent implementation [196]. In response, the NIS2 Directive was adopted in 2022 to expand the scope of the original directive by enforcing additional entities and sectors to improve their cyber security capabilities [270]. At a national level, The United Kingdom (UK) created the National Cyber Security Centre (NCSC) in 2016 to provide advice and support to public and private sectors alike [186]. As part of the NIS Directive's adoption, the NCSC created the Cyber Assessment Framework (CAF) to provide British organisations and CNI with guidance for improving their cyber security [185]. Objective D, in particular, aims to aid stakeholders in improving their R&R capabilities, including planning for cyber incidents and lessons learnt from these.

Bridewell published a report in 2022 outlining trends to expect for 2023 concerning the cyber security of CNI and ICS/OT [25]. Firstly, while nation-state-sponsored attacks are still a significant threat, the rise of "cyber attacks as a service" has led to a significant increase in cyber attacks originating from criminal organisations, shifting away from traditional forms of illegal activities towards cyber crime. Due to the cost of living crisis that has affected the UK and other countries across the globe since 2021, many criminal organisations are starting to target vulnerable insiders within organisations as a means of easily gaining initial access to environments or sensitive information. In parallel, CNI organisations are also predicted to reduce their cyber security budget, despite the NIS2 Directive, due to the costs and pressures caused by the cost of living crisis, leading to a higher probability of successful attacks against these. Finally, due to the technological "arms race" that organisations face against threat actors, security implementations and tools are becoming increasingly complex and varied, leading to reduced visibility and overall reduced security capability.

Overall, while guidance supporting the evolution of ICS/OT security, and compliance with regulations, equips operators with a starting point, it may not be complete and actionable. In particular, with the expected trends for cyber attacks in 2023, ensuring that cyber security capabilities are efficiently implemented is paramount. This is of particular interest for cyber incident R&R best practice, where CNI and its supporting ICS are used for critical services.

1.2 Research Questions

The problem space identified from section 1.1 can be summarised as such:

- Cyber security for ICS/OT presents additional challenges than traditional IT.

- The convergence of IT and OT through Industry 4.0 has led to an increased attack surface.
- Because of this, threat actors have increasingly begun to target CNI.
- Cyber incident R&R is an essential part of the cyber security lifecycle but challenges exist for effectively planning against the observed rise in cyber attacks targeting CNI and ICS/OT.

Research questions (RQs) must be developed before research can be organised and focused on the problem space outlined above. These research questions are are as follows:

- **RQ1**: How has the convergence of IT and OT affected the way that preparation for ICS/OT cyber incident response and recovery needs to be handled?
- **RQ2**: How effective are current ICS/OT cyber incident response and recovery capabilities in practice?
- **RQ3**: Which areas of ICS/OT cyber incident R&R are significantly lacking?
  - **RQ3.1**: How can these areas be improved to better prepare for cyber attacks targeting ICS/OT?
  - **RQ3.2**: Could an approach be developed in these areas to improve cyber incident R&R?

By answering these questions, this thesis contributes to academic and industry knowledge through an in-depth exploration of cyber incident R&R for ICS/OT. By identifying gaps in literature and current practices, an approach can be identified and developed for improving cyber incident R&R for ICS/OT.

1.3 Thesis Structure

This chapter discussed the role of ICS/OT within CNI and the effect that industry 4.0 has had on how cyber security for ICS/OT is handled, especially for cyber incident R&R. In order to answer the research questions derived from this problem space, the thesis is structured as follows:

Chapter 2

Chapter 2 provides background research to understand the importance of cyber incident R&R for ICS/OT. In this chapter, an analysis is firstly made of cyber attacks targeting CNI and ICS/OT, accentuating the need to effectively prepare for responding to and recovering from these. Secondly, an analysis between IT and OT is conducted to better understand why traditional cyber security for IT is not always applicable to ICS/OT environments. Through this, Chapter 2 answers RQ1.

Chapter 3

Chapter 3 provides an analysis of current R&R practices for ICS/OT through an analysis of existing standards and guidelines, and engagement with stakeholders. Through this analysis, gaps in current practices can be identified for the development of novel approaches to improving cyber incident R&R. This chapter, therefore, enriches the answers to RQ1 and answers research questions RQ2 and RQ3.

Chapter 4

Chapter 4 proposes a framework for aiding stakeholders in assessing and improving their cyber incident R&R capabilities through standards and guidelines. Through the development of this framework, and the outputs of Chapter 4, specific areas of R&R for ICS/OT that are significantly lacking are identified for further exploration; specifically adversary-centric security testing, an assurance technique used during risk management for preparation of cyber incidents. As with Chapter 3, this chapter contributes to answering RQ3.

Chapter 5

Chapter 5 provides an analysis of adversary-centric security testing for ICS/OT and how the challenges that exist within this area contribute to the difficulties encountered in Chapters 3 and 4. In doing this, this chapter identifies areas for development to improve cyber incident R&R through adversary-centric security testing. Therefore, Chapter 5 provides more depth for answering RQ1 and RQ3, and provides direction for answering RQ3.1 and RQ3.2.

Chapter 6

Chapter 6 leverages the outcomes of previous chapters to propose a framework that is intended to improve cyber incident R&R for ICS/OT through the scoping of adversary-centric security tests. The framework provides methodologies for factoring safety and operational risk into

the overall scoping process of these engagements. The result of the framework proposed in this chapter is the actualisation of RQ3.1 and RQ3.2.

Chapter 7

Chapter 7 presents a qualitative study with experts involved in adversary-centric security tests for ICS/OT environments as part of an evaluation of the scoping framework proposed in the previous chapter. In this study, the framework is presented to participants alongside an example application scenario for it to describe how the framework could be used to enable adversary-centric security tests on ICS/OT that take into consideration existing safety and operational risk. In doing so, the study evaluates the framework's accuracy, reliability, validity, and applicability in practice through the thoughts and opinions of these experts. This evaluative study confirms that the proposed framework in Chapter 6 is suitable for answering RQ3.1 and RQ3.2.

Chapter 8

Chapter 8 concludes the thesis by reflecting on the research questions and how these were answered. In this chapter, the contributions throughout the thesis are summarised and future work is proposed to enhance the contributions developed within the thesis.

Chapter 2

Background

Chapter 1 introduced the concept of ICS/OT and the importance of these systems for CNI. Due to the convergence of IT and OT, through Industry 4.0, securing these systems against threats has become increasingly complex. In recent years, there has been a global push to improve cyber security capabilities within organisations, primarily in response to the dramatic increase in targeted cyber attacks [171]. Increasingly, such attacks have started targeting networks of CNI, which, as a reminder, are systems that are essential for the smooth operation of a country's economy and society [205]. Successful cyber attacks on CNI can have serious consequences, as observed in the Stuxnet attack of 2010, which caused significant delays for the development of the Iranian Nuclear Programme [199]. Such attacks can also bring danger to civilian life by affecting critical environments such as electrical grids, emergency services and transportation services.

This chapter, firstly, emphasises the need to secure ICS/OT and CNI against threats through an analysis of cyber attacks targeting these environments and the subsequent trends that can be learnt from these attacks. Secondly, this chapter aims to answer RQ1 by demonstrating how the differences between IT and OT create additional challenges for securing industrial networks in the era of Industry 4.0.

2.1 Cyber Attack Trends on Industrial Control Systems

As described in Chapter 1, the digitisation of the industrial process through Industry 4.0 has led to industrial environments' attack surface increasing dramatically in recent years. Because of this, several cyber attacks targeting CNI and ICS/OT have been observed. Such attacks can have serious consequences, including disrupting critical infrastructure and industrial processes, causing financial losses, and even endangering human lives. The aim of this

section is to analyse the trends relating to these attacks and discuss the potential risks and challenges that these pose.

2.1.1 Overview of Attacks

Table 2.1 provides a summary of 43 attacks from 1988 to 2021 targeting ICS/OT that were selected for analysis. Information on these attacks was collected through public data feeds such as news sources, white papers and case studies. These attacks were selected based on whether observed impact was observed within levels 0 to 3 of the Purdue Model (see Figure 1.1). This means that while some of these attacks may not have originally targeted ICS/OT, because of cascading effects between the IT and OT networks, the operational process was still impacted. The following information is provided by Table 2.1:

- Attack: common name of the attack.
- Date: The year the attack was observed.
- Initial Access: The technique used to gain initial access.
- Threat Actor: The type of threat actor responsible for the attack.
- Sector: the sector that was affected by the attack.
- Impact: The impact caused by the attack.
- References: Sources used to collect information on the attack.

Despite the fact that the majority of ICS/OT were not connected to the internet in the 1980s and early 1990s, the extrapolated data from these attacks was found to be critical for the analysis. This additional information aids in identifying a more accurate transition over time of the attack vectors, sectors, and impact, resulting in a more accurate analysis. Because of the length of time since these attacks were recorded, this data should be interpreted with caution due to the potential for inaccuracy of public information. To provide as much accuracy as possible, these attacks were cross-referenced with multiple sources (RISI [27], ICS/OT related news sources, white papers, and case studies). The Siberian Pipeline Explosion (1982) [102] is an example of exclusion from this process, as there is uncertainty around the existence of this attack.

Six of the eight attacks prior to the year 2000 are caused by insiders. Despite a shift away from this attack type, at least in terms of publicly available information, documents are still being published that identify these threats as an undefeated problem within the domain of

ICS/OT [7, 204, 25]. As a result, valid information on old attacks can still provide cyber security practitioners with a better understanding of the evolution of attack vectors, threat actors, impact, and targeted sectors and locations.

2.1.2 Attack Trends and Lessons Learnt

Methodology

Using STIX (Structured Threat Information eXpression) Objects [208] and the ATT&CK ICS Framework Tactics [174, 178], the extracted information from the attacks summarised in Table 2.1 are presented in Figure 2.1. STIX serves as a standardised language that aims at providing comprehensive Threat Intelligence data in a structured way [208]. The following STIX Objects have been used in the data extraction: Campaign, Course of Action, Identity, Indicator, Infrastructure, Intrusion Set, Location, Malware, Observed Data, Threat Actor, Tool, and Vulnerability. Additional technical information from the MITRE ATT&CK framework for ICSs was also leveraged [178] covering Initial Access and Impact. This framework is described as "a knowledge base useful for describing the actions an adversary may take while operating within an ICS/OT network. The knowledge base can be used to better characterise and describe post-compromise adversary behaviour." [173].

Four key categories were identified during the extraction of attack information: Threat Actors, Initial Access Techniques, Impact, and Targeted Infrastructures & Locations. These categories serve as the basis for analysis. Each of these categories is supported by using existing frameworks and/or taxonomies; justifying their use within this analysis.

When analysing the data corresponding to Threat Actor information, the following STIX objects were used: Identity, Intrusion Set, and Threat Actor. Additionally, the Threat Actor Taxonomy provided by the Center for Internet Security (CIS) [37] was also chosen to categorise different threat actor groups. Although many taxonomies exist for classifying threat actors such as the one provided by the Cybersecurity and Infrastructure Security Agency (CISA) [49], the CIS taxonomy was selected as it provides a clear distinction between Threat Actors based on their knowledge, skills, abilities, motivations, and resources. The selected taxonomy is as follows:

- Nation State or Nation State Sponsored are groups that may be part of a nation state government branch or are provided resources and funding from a nation state. They often have immense resources and funding for carrying out their mission, and their motivations are often political, military, or to conduct espionage.

Attack	Date	Initial Access	Threat Actor	Sector	Impact	References
PLC Password Change	1988	Workstation Compromise	Insider	Manufacturing	Denial of Control	[27]
Ignalina Nuclear Power Plant	1992	Workstation Compromise	Insider	Civil Nuclear	Loss of Productivity and Revenue	[17, 27]
Chevron Refinery Emergency Alarm System	1992	Workstation Compromise	Individual	Chemical	Loss of Productivity and Revenue	[52]
Salt River Project	1994	Internet Accessible Device	Individual	Energy and Water	Loss of Productivity and Revenue, Disk Wipe	[27]
Omega Engineering	1996	Workstation Compromise	Individual	Manufacturing	Disk Wipe	[27]
Worcester, MA Airport	1997	Internet Accessible Device	Individual	Transport	Loss of Productivity, Revenue, Availability, and Safety	[27]
Gazprom	1999	Unknown	Organised Group + Employee	Chemical and Energy	Loss of Productivity and Revenue	[27, 57]
Bradwell Nuclear Power Plant	1999	Workstation Compromise	Insider	Civil Nuclear	Disk Wipe	[163]
Maroochy Water System	2000	Wireless Compromise	Insider	Water	Damage to Property	[27, 57]
Cal-ISO System	2001	Unknown	Nation State	Energy	None Disclosed	[27, 182]
Virus on Manufacturing System	2001	Spearphishing	Nation State	Manufacturing	Loss of Productivity and Revenue	[27]
Houston Port	2001	Internet Accessible Device	Individual	Transport	Loss of Productivity and Revenue	[27, 156]
Gas Processing Plant	2001	Trusted Relationship	Supplier	Chemical	Loss of Productivity and Revenue	[27, 245]
PDVSA	2002	Internet Accessible Device	Organised Group	Chemical	Loss of Productivity and Revenue, Disk Wipe	[27, 245]
Flight Planning Computer	2003	Unknown	Individual	Transport	Loss of Productivity and Revenue	[244, 27]
CSX Train Signalling System	2003	Spearphishing	Unknown	Transport	Loss of Productivity and Revenue	[27]
Contractor Infects SCADA Network	2004	Replication Through Removable Media	Unknown	Food	Loss of Productivity and Revenue	[27]
Daimler Chrysler Plants	2005	External Remote Service	Individual	Manufacturing	Loss of Productivity and Revenue	[57]
Tehama-Colusa Canal	2007	Workstation Compromise	Individual	Water	Damage to Property	[27]
Lodz Tram System Hacked	2008	External Remote Service	Individual	Transport	Loss of Safety	[135, 35]
US Power Grid	2009	Internet Accessible Device	Nation State	Energy	None Disclosed	[198]
Hospital HVAC	2009	Workstation Compromise	Insider	Health	Loss of Safety	[27]
Night Dragon	2009	Exploit Public-Facing Application	Organised Group	Energy	Theft of Operational Data	[166, 57, 198]
Sality Virus Infects DVS Servers	2009	Unknown	Unknown	Chemical	Loss of View	[27, 169]
Stuxnet	2010	Replication Through Removable Media	Nation State	Civil Nuclear	Damage to Property, Manipulation of View and Control	[157, 153, 170, 27]
Shionogi	2011	Workstation Compromise	Individual	Health	Disk Wipe	[27]
Niagra AX	2012	Internet Accessible Device	Unknown	Manufacturing	Manipulation of Control	[274, 27]
Espionage on Iranian CI	2012	Replication Through Removable Media	Nation State	Chemical	Theft of Operational Data, Unintentional Disk Wipe	[170, 286, 21]
Turbine Control System	2012	Replication Through Removable Media	Organised Group	Energy	Loss of Productivity and Revenue, Theft of Operational Data	[27, 243]
Rye Brook Dam	2013	Internet Accessible Device	Organised Group	Water and Energy	None Disclosed	[100]
European Public Utility Services Attacked	2014	Spearphishing	Organised Group	Various	Denial of Service, Theft of Operational Data	[112]
German Steel Mill	2014	Spearphishing	Unknown	Manufacturing	Damage to Property	[27, 154]
Ukrainian Energy	2015	Spearphishing	Organised Group	Energy	Loss of Productivity and Revenue	[287]
Ukrainian Energy	2016	Spearphishing	Organised Group	Energy	Disk Wipe, Loss of Productivity and Revenue, Loss of Safety	[62, 57]
Wolf Creek	2017	Spearphishing	Organised Group	Civil Nuclear	None Disclosed	[221]
Cadbury Factory Attack	2017	External Remote Service	Organised Group	Food	Loss of Productivity and Revenue	[271, 48, 96]
Triton/Petro Rabigh	2017	Workstation Compromise	Nation State	Chemical	Denial of Control, Loss of Safety	[136, 87]
Norsk Hydro	2019	Spearphishing	Unknown	Manufacturing and Energy	Loss of View	[272, 254]
Triton/Undisclosed	2019	Workstation Compromise	Nation State	Undisclosed	Denial of Control, Damage to Property, Loss of Safety	[240, 136]
Hackers Target Oil Producers	2020	Spearphishing	Unknown	Chemical	Theft of Operational Data	[14, 73]
Israeli Water Facilities Attacked	2020	Internet Accessible Device	Organised Group	Water	None Disclosed	[149, 150]
Cyber-Attack on Shahid Rajaie Port	2020	Unknown	Nation State	Transport	Loss of Productivity and Revenue	[134, 228]
Honda Factories Cyber Attack	2020	Spearphishing	Unknown	Manufacturing	Denial of Control	[61, 44, 164]

Table 2.1 Summary of Attacks

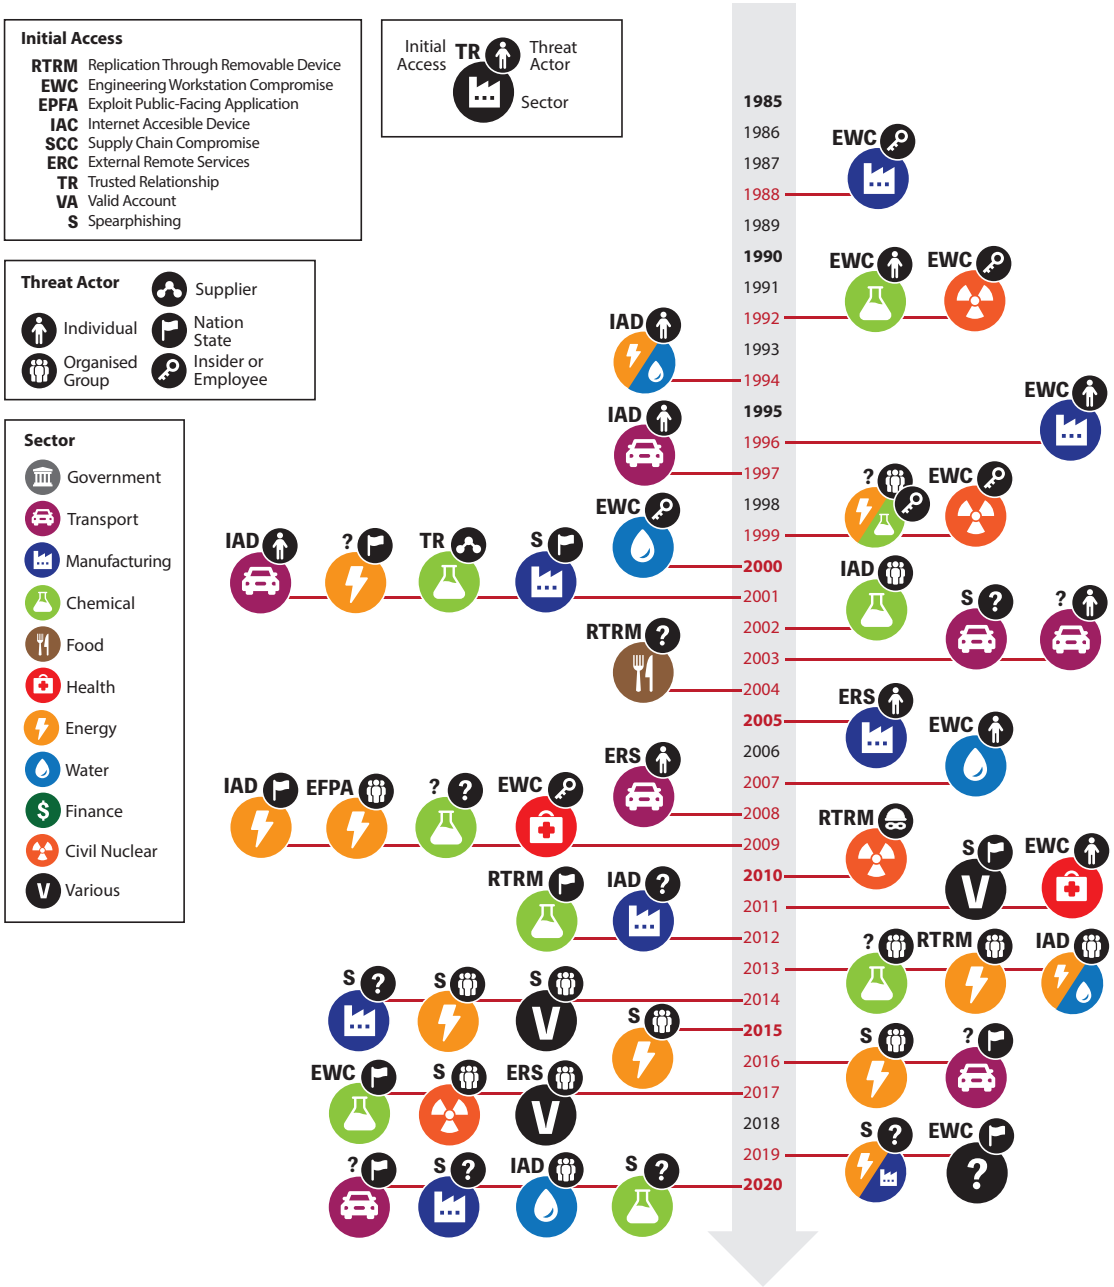

Fig. 2.1 Timeline of Attacks against ICSs [171]

- Organised Groups include groups of cyber criminals or cyber terrorists with average to considerable resources for carrying out attacks. Their motivations are ideological, financial, or social.
- External Individuals concern individual actors with no prior access to the systems they wish to exploit. They often have little resources, and their motivations are financial or personal.
- Insiders are trusted individuals within an organisation that already have some access to the systems they intend on exploiting (often in part to being an employee within the organisation).

For the analysis of Initial Access techniques, identified techniques were mapped to the tactics and techniques from both the MITRE ATT&CK Framework [174] and the MITRE ATT&CK for ICS Framework [178]. These frameworks have been selected as some techniques found in the IT-specific framework, such as the use of Valid Accounts, were also applicable within an industrial context. Similarly to this, each attack's impact was also categorised following both these frameworks.

Identified infrastructures were categorised based on the NPSA's taxonomy of National Infrastructure Sectors [205]. Although not a national infrastructure, the manufacturing sector has also been considered within the analysis as this sector often involves the use of ICS/OT networks and is the target of several examined attacks.

The following subsections expand on Figure 2.1 to identify associated trends that have emerged throughout the history of ICS/OT attacks, and what lessons can be learnt to prepare for potential future attacks. While an extensive set of resources has been used when extracting data on the ICS/OT cyber attacks, some attacks do not have sufficient access to resources to confidently identify the information required for a comprehensive analysis. This is most often due to information on specific attacks being classified or unavailable.

Threat Actor

One of the most critical components in current threat intelligence is understanding threat actors, their behaviour, motivations, and capabilities. For this section of the analysis, the Identity, Intrusion Set, and Threat Actor STIX Objects were used alongside the CIS Threat Actor Taxonomy.

Prior to 2009, most attacks (13 out of 20) are confirmed to have been conducted by individuals, both external and internal. From 2009, a clear transition to larger and more organised groups can be observed. 15 of 23 attacks are either confirmed or allegedly from

nation state-sponsored groups or organised groups. When analysing trends from threat actors, it is also important to note the motivation behind these attacks. Many individuals carried out attacks due to personal reasons for either financial gain or as methods of retribution. This can be seen in attacks such as the Bradwell Nuclear Power Plant attack of 1999 [163] or the Houston Port attack of 2001[27, 156]. However, Organised Groups' motivations were mostly political with the aim of conducting espionage or disruption as observed in the Stuxnet attack of 2010 [199, 157, 27].

Two clear trends concerning threat actors have been identified in the observed attacks over the past 32 years. While an increase in complexity of systems and an increase in security awareness suggests that it is more difficult for single individuals to carry out attacks due to limited skill and resources, there is also a noticeable increase in organised threat capability, often provided with extensive resources through nation-state funding. Although basic security strategies are commonly being implemented, resulting in fewer incidents from simple attack vectors such as poor access control or common vulnerabilities, the introduction of methods for increasing interconnectivity and the complexity of modern systems has increased the possible attack surface for groups with considerable resources to discover and exploit.

To mitigate security risks from individuals, practitioners must ensure the implementation of fundamental security strategies within their organisation. This includes but is not limited to the following: resilient access control such as revoking credential access from terminated employees or ensuring that only authorised members have access to critical systems, and thorough vetting of personnel. A plethora of existing standards and guidelines such as the NIST SP 800 series [201] or the IEC 62443 series [118] can be consulted for practitioners to assess their current security strategies and reevaluate them if necessary. Despite acting as individuals, insiders present additional security risks to organisations due to their already possible access to critical assets and their knowledge of the intricacies of the organisation they are employed by. For this reason, many governmental and standard bodies provide specialised guidance for these threats such as the resources provided by CISA which include methodologies for appropriately identifying and responding to insider threats [50]. Such recommendations include implementing rigorous vetting when hiring new employees, detecting changes in emotional behaviour due to psychological factors, and more. Similarly, academic articles can also provide information on mitigating risks caused by insider threats such as the survey conducted by Homoliak et al. [105] on insider threat taxonomies, analysis, modeling, and countermeasures. Outputs from this survey include mitigation and prevention recommendations such as decoy-based, opportunity-based or anomaly-based detection methods.

To this day, organised groups constitute the most considerable risk to critical infrastructures. To combat this growing threat, countries across the globe have adopted the use of national cyber security organisations such as the National Cyber Security Centre in the United Kingdom [190] or The Cybersecurity and Infrastructure Security Agency in the United States [39]. These serve as central points of information and guidance for organisations from private and public sectors alike. Practitioners are highly recommended to ensure that they make regular use of the guidance and threat intelligence provided by these to improve their cyber security and incident response capabilities. A push has also been observed for organisations to share Threat Intelligence through less centralised methods such as open-source Threat Intelligence feeds like Proofpoint's Emerging Threats Intelligence software [224] or the FBI's InfraGard, which is specifically tailored towards Critical Infrastructures [76].

Initial Access

The Initial Access techniques from the MITRE ATT&CK and ATT&CK ICS Frameworks [174, 178] were selected when categorising the techniques identified from each attack within Table 2.1.

Abuse and utilisation of a valid account through the compromise of an engineering workstation (ATT&CK ID T1078 and T0818) have been identified as the most commonly used techniques to gain a foothold into target systems throughout the first half of the investigated time period. This trend suggests that early attacks on ICSs relied heavily on the abuse of an existing level of trust and access. This is most likely because ICS/OT networks were traditionally disconnected from any other networks and used proprietary protocols. Therefore, either physical security needed to be bypassed or an already existing level of access is required (e.g. an employee). An example of bypassing physical security can be observed in the Lodz Tram System attack of 2008 [135, 35].

Unlike the first half of the investigated time period, a wider variety of Initial Access techniques are used in the second half. These techniques include exploitation of External Remote Services (ATT&CK ID T0822), access through Internet Accessible Device (ATT&CK ID T0883), Replication Through Removable Media (ATT&CK ID T0847), and use of Spearphishing Attachment (ATT&CK ID T0865). From 2013, there is a noticeable shift from using technical Initial Access techniques towards the use of social engineering methods such as spear-phishing. Examples of these can be observed in the German Steel Mill attack of 2014 [154, 27], the Norsk Hydro attack of 2019 [272, 254], and the Honda factory attack of 2020 [61, 44, 164].

These identified trends suggest an evolution of the importance allocated towards ICS/OT cyber security over the years. Historically, ICSs were mostly protected at the network level

through the use of air-gapping and proprietary protocols, making it extremely difficult for external actors to gain access to these systems [28]. However, with the advances in modern technology and the integration of standardised protocols within industrial networks such as TCP/IP, the attack surface has increased. As technical security strategies have improved over the years, most recent attacks have turned to rely on some form of social engineering or human error to gain an initial foothold into targeted systems. This highlights the importance of both providing cyber security training to all employees within an organisation and ensuring that organisations correctly implement a robust security culture within work environments. Practitioners are advised to consult their national cyber security organisations' guidance regarding minimising attack surface and social engineering awareness. To highlight the importance of social engineering awareness, the NCSC in the UK has currently published a total of 58 guidance resources on phishing exclusively [191]. While training is important, providing this alone does not provide adequate protection against social engineering attacks. Organisations should also ensure that a resilient security culture is implemented within work environments. This includes preventing risky behaviour by establishing stress free environments to minimise mistakes (e.g. accidentally opening an email attachment due to lack of attention or holding the door to a restricted area open to someone without first checking access privileges). Guidance on establishing a robust security culture can be found through various sources such as the open source security culture framework [230].

Although initial access into the target system is commonly carried out through social engineering, follow up tactics such as Lateral Movement, Data Collection, or Command & Control are still often executed using either zero-day exploits or known vulnerabilities. Therefore it is also recommended for practitioners to keep abreast of recent Common Vulnerability and Exposures (CVEs) through sources such as MITRE's CVE Database [175] or the National Institute of Standards and Technology (NIST) National Vulnerability Database (NVD) [203]. Keeping up-to-date with these vulnerabilities alone is, however, not sufficient enough to confidently prepare for cyber attacks. Making use of Assurance Techniques such as Document Reviewing or Testing also provides benefits towards an organisation's cyber security capabilities and should, therefore, also be considered [148].

Impact

Similarly to Initial Access Techniques, each attack's impact was also categorised following the MITRE ATT&CK ICS Framework [178].

Unlike the trends identified for initial access techniques, there is no distinct shift in observed impact on systems. This is most likely due to the impact of attacks being closely linked to each attack's motivation rather than an evolution in adversary capabilities. Therefore,

identified impacts have been associated with the three following attack motivations: Financial Gain, Espionage/Information Gathering and Disruption/Sabotage.

Only 2 of the 43 attacks were conducted due to financial motivation. The impact from these attacks includes the loss of view or control of systems (ATT&CK ID T0829 and T0813) which often also resulted in a Loss of Productivity and Revenue (ATT&CK ID T0828). This occurs mostly due to the type of malware used in these attacks: ransomware; resulting in the encryption or removal of essential files required for operation. The low frequency of these attack types suggests that financially motivated attacks target primarily IT systems. This is partly due to the ratio of IT to industrial systems (more systems to attack results in a higher possible monetary gain). At least 15 ransomware attacks have targeted non-industrial organisations such as Universities or law firms in the first half of 2020 alone [46]. Although financially motivated attacks target IT systems more than industrial systems, poor network architecture management can result in the spread of malware from IT systems to industrial systems. This was observed in the Air Canada attack of 2003, where the Blaster Worm, a malware targeting Microsoft Windows initially, spread into the air company's flight planning network [244, 27]. This demonstrates the importance of correctly segregating IT networks from ICS/OT networks to prevent IT incidents from affecting ICSs as proposed with the extended Purdue Enterprise Reference Architecture and the use of a Demilitarized Zone to separate IT and industrial networks from each other, for example [40].

8 of the 43 attacks were conducted in order to steal information. In most cases, this resulted in the successful theft of operational data (ATT&CK ID T0882) such as building blueprints, network topologies, confidential documents, or user credentials. Additionally, some of these attacks caused system disruption with disk wipes (ATT&CK ID T1561) or system crashes resulting in a Loss of Productivity and Revenue (ATT&CK ID T0828). This can be observed with the Flame malware deployment, used to conduct espionage in Middle Eastern countries and cause disk wipes [170, 286, 21]. On many occasions, these attacks exist as precursors to activities with the intent to cause disruption; often cyber-related, but not always. If system operators manage to detect an attack that has resulted in the theft of confidential or valuable information, they should be prepared for a follow-up attack with a potentially more disruptive goal.

The majority of the identified attacks were conducted to cause sabotage or disruption whether it be by a disgruntled ex-employee targeting a specific organisation as part of a vengeance ploy, or by a nation-state targeting systems that could damage another country's economy or operations. This often resulted in a Denial, Manipulation or Loss of Control and/or View (ATT&CK ID T0813, T0831, T0827, T0815, T0832, and T0829). Consequently, a Loss of Productivity and Revenue (ATT&CK ID T0828) was also observed with Loss

of Safety and/or Damage to Property (ATT&CK ID T0880 and T0879) in some cases as seen in the attack on the German Steel Mill in 2014 which caused damage to a furnace due to it being unable to shutdown [27, 154]. This highlights the importance of effectively securing ICSs and responding effectively to cyber incidents if prevention techniques fail. Compared to traditional IT systems, attacks on industrial systems also can cause a loss of safety and therefore, a danger to life. Therefore, practitioners should ensure that their organisation's response plans are thorough and make good use of existing guidance and standards available [263].

Infrastructure and Location

To conclude the analysis, each attack's targeted infrastructure and location were explored using the associated STIX Objects. This was done to determine if any specific infrastructures were more commonly targeted and if there were any changes in targeted infrastructures over the years. Identified infrastructures were categorised based on the NPSA's taxonomy of National Infrastructure Sectors [205].

Prior to 2009, a variety of targeted infrastructure and locations can be observed. This is partly due to the threat actors involved behind these attacks: as individuals from specific organisations were behind most of the incidents, no infrastructure or location-specific trends were identified. This can be seen, for example, in the Texas Hospital HVAC attack in 2009 [27] where a security guard working at the hospital was responsible for the attack. The corresponding infrastructure was targeted not because it was a hospital but because the attacker was employed there.

From 2010 onward, a clear shift towards targeting the chemical and energy sector can be observed. This is because of the two following reasons: the associated impact caused by targeting these sectors and the motivation behind these attacks. The destructive impact associated with the disruption of the energy sector could cause detrimental consequences to a broader array of infrastructures that require electricity to function. In contrast, an attack on infrastructure, such as oil refinery plants within the chemical sector, could have a severe economic impact on the associated nation. This is especially true for Middle Eastern countries such as Saudi Arabia where the petroleum sector accounts for 42% of the country's GDP [6]. Therefore attacks such as the one targeting Petro Rabigh in 2017 [136, 87] not only had the potential to cause a danger to life but could also have had severe consequences on the country's economy. Attacks such as the 2015 attack on the Ukrainian Energy Sector which caused power outages for over 80,000 residents [287] have the potential to affect other infrastructures, taking systems such as assembly lines, life-saving hospital apparatus, or

chemical processing machines offline. It can be inferred that these sectors have been targeted because of the impact these attacks can have on other, energy-dependent, sectors.

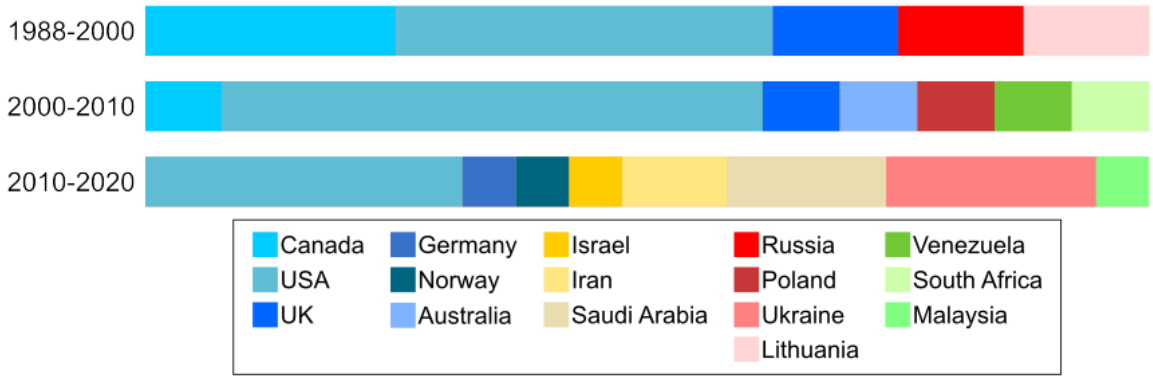

Fig. 2.2 Evolution of Cyber Attack Locations [171]

The location of these attacks could suggest an evolution of both post-Soviet and Middle Eastern conflicts towards a more digital environment. Many western countries are allegedly behind attacks targeting post-Soviet or Middle Eastern countries and vice versa. This cyber warfare can be observed through the back and forth attacks between Iran and Israel, for example [149, 150, 134, 228]. This evolution is illustrated in Figure 2.2. The risk-reward aspect of conducting cyber warfare over the use of a physical medium must also be considered. The little risk and high reward of disrupting a nation-state through a cyber attack lead to the logical increase in these methods' use over more traditional ones such as physical interventions or economic sanctions.

2.1.3 Summary

As discussed in the previous sections, multiple trends surrounding past attacks on ICSs were identified. These include Threat Actors behind the attacks, Initial Access techniques to gain a foothold into the target systems, the impact and motivation of each attack, and the attacks' target infrastructure and location.

Prior to 2009, the Threat Actors responsible for the examined attacks were mostly individuals, whether external or internal, targeting infrastructures where these were employed. Initial Access techniques, therefore, involved the use of an already existing level of access. The motivation behind each attack was mostly personal, targeting organisations as retribution, and accordingly, the impact of these attacks was mostly disruptive in nature. During this time period, there was not a large emphasis on ICS/OT security as these systems mainly were protected through physical means such as segregated networks or the use of propriety

technology. Weak cyber security policies such as Access Control enabled employees such as security guards to access systems on an industrial network without much difficulty. Although organisations have accorded much greater importance towards cyber security in recent years, it is still essential that stakeholders ensure their security policies are implemented thoroughly throughout the organisation, and make good use of existing guidance and guidelines.

From 2009 onward, a shift can be observed towards attacks conducted by more organised groups such as cyber criminal or nation state-funded groups. Therefore, the motivation behind these attacks also shifted towards more political reasons such as espionage or sabotage and targeted more critical infrastructure such as the energy sector. Nation-states with a long history of tension have adopted cyber space as an additional battleground against each other. As security awareness has increased considerably in the past decade, threat actors have also become more reliant on exploiting human vulnerabilities with social engineering when gaining an initial foothold into target systems. Due to this ever-expanding threat landscape that we now face, organisations must work together to understand and prepare against such threats effectively. Keeping in regular contact with national cyber security organisations and sharing threat intelligence between organisations have become essential.

There is a clear separation between the identified trends before and after 2009. This shift in trend coincides with the public exposure of the Stuxnet attack of 2010. This attack was described as the first known use of malware that was crafted to target ICSs specifically and is also the first known use of a cyber weapon [74]. Because of the detrimental effects that Stuxnet had on Iran's nuclear program, this attack was highly publicised and discussed throughout the security community and media alike. It highlighted the importance of defending ICSs against malicious actors as it was now known how damaging these types of attacks could be [255]. While stakeholders were made aware of ICS/OT security's importance, malicious actors were also exposed to the possibilities of executing a cyber attack on ICS [53]. This newfound interest in ICS/OT environments from attackers, coupled with the shift in ICS/OT environment construct (e.g. broader inter-connectivity), converged during the middle of the investigated time period. It can therefore be considered that both of these factors contributed to the apparent change in trends that has been observed from 2009 onward.

2.1.4 Attacks Post-2021

While the above analysis covers a considerable period, it does not include attacks that have occurred after 2021. These are worth discussing; however, they have been excluded from the presented study due to a lack of available and reliable data, in part because of the recency of these attacks.

Many organisations, including Kaspersky [140], Fortinet [78], and Dragos [63], released their annual report on the state of ICS/OT cyber security at the end of 2022. Within these, several cyber attacks targeting industrial sectors were reported. For example, in June 2022, multiple steel production companies in Iran were found to have had the vacuum degassing machinery of their mills compromised, leading to a loss of structural integrity and increase in defects within the produced steel [257]. In August 2022, South Staffordshire Water in the UK was the target of a ransomware attack on their corporate IT systems [180]. While the enterprise network was disrupted, the company stated that the attack did not affect its "ability to supply water" to its customers. Additionally, despite the Russian invasion of Ukraine being largely conventional to date, The UK Office for Budget Responsibility [211] reported that several cyber attacks targeting Ukrainian CNI, including energy, nuclear, communications and defence, had been observed.

While many more attacks are presented within these reports, a common reoccurrence that was identified was that these attacks, while targeting industrial sectors such as manufacturing or water, did not, in most cases, cause impact to the operational process. This phenomenon has recently been observed by Derbyshire [55], who noted that most recent attacks that have been reported as attacks on OT are, in fact, "OT-but-not-OT" attacks due to the majority of these not leading to operational impact or, where OT impact was observed, this being due to cascading effects from compromise of the IT network. However, despite this emerging trend, it is crucial to acknowledge these attacks' impact on OT environments, regardless of their original intent. Therefore, preparing for these attacks to minimise their impact remains essential.

2.2 Comparison of Operational and Information Technology

Despite the urgency to effectively defend CNI against cyber attacks, as demonstrated in Section 2.1, the fundamental differences between OT systems, often used in CNIs, and the more traditional Information Technology (IT) systems, makes it difficult to transfer skills and techniques from one domain to the other [184], especially with their convergence due to Industry 4.0. As opposed to IT systems, which prioritise the handling of data or information, OT systems are used to ensure the successful operation of operational processes. The design of these two technology-types needs to, therefore, take into consideration these differences which also affects the way that cyber security measures are implemented.

This section aims to compare and contrast IT and OT, highlighting their key differences and similarities, and exploring the challenges associated with their integration within the context of cyber security.

2.2.1 Methodology

For this analysis, standards and guidelines have been selected which discuss asset management in detail. The scope of the sections discussing asset management within these articles is to guide stakeholders with asset management and concentrate on activities such as implementing a new asset within a network, transferring an existing asset to another network, and continuous hardware and software monitoring. As presented in these documents, asset management is considered a critical part of an organisation's cyber security lifecycle as mismanagement of this process can adversely affect other phases or categories such as resilience or incident response. By selecting the categories discussed during asset management sections, distinct criteria can be identified for comparing OT and IT systems. While several publications on asset management specifically were identified such as NIST SP 1800-5, which discusses IT Asset Management for large financial services organisations [188], most of these articles referenced the following two documents: The NIST Framework for Improving Critical Infrastructure Cybersecurity [187] and the ISO/IEC 27000-series [129, 130]. The sections from these documents detailing asset management were therefore selected as the basis for this comparison.

The first document that was selected for the comparison is the NIST Framework for Improving Critical Infrastructure Cybersecurity [187]. This document was created following the Cybersecurity Enhancement act of 2014, which tasked NIST with developing a cyber security risk framework for operators and owners of Critical Infrastructures [278]. The framework scope is to provide guidance to operators and owners of Critical Infrastructure to improve their cyber security activities and help implement resilient cyber security strategies within an organisation's larger risk management process. A large portion of this framework provides guidance on effectively managing assets of critical infrastructures as part of the cyber security lifecycle. As ICSs are often found within critical infrastructure networks, this document can mostly be extended to include any operators or owners of OT assets.

The second document used for this comparison is the ISO/IEC 27000-series. This series of standards, mainly ISO/IEC 27001 and 27002, provide best practice guidance for information security management [129, 130]. Having a broad scope, these standards can be used by all types of organisations. While generally IT-focused, the sections of these standards that do not go into technical detail can still mostly be used by operators and owners of ICSs.

Section Title	ISO/IEC 27001/2	NIST Framework
Asset Management	A.8	ID.AM
Physical Security	A.11	PR.AC
Software Management	A.12.5	ID.AM
Communications Security	A.13	ID.AM, PR.AC, PR.DS
Security Policies	A.5	ID.GV
Awareness and Training	A.7.2.2	PR.AT

Table 2.2 Equivalent Sections from the NIST Framework and ISO/IEC 27001/2

However, it is recommended that more OT-focused standards such as ISO/IEC 27019 be consulted as well, especially for the energy industry [131].

To begin with the comparison, all relevant sections from the NIST Framework for Improving CI Cybersecurity and ISO/IEC 27001/2 on Asset Management were extracted. These sections allowed for categorisation of different aspects of IT and OT systems for the comparison. However, despite these two documents catering towards OT and IT, respectively, the topics related to Asset Management are identical in content, as shown in Table 2.2. From this, the following categories were extracted to use in the analysis:

- Hardware: function, manufacturing etc.
- Software: underlying programming, patching etc.
- Network: network topology, communication protocols etc.
- Socio-Technical: governance, policies, education, training etc.

2.2.2 Hardware Characteristics

One of the most fundamental differences between OT and IT-based systems is their hardware. This is due to, essentially, their function within their respective environments. As per the name, Information Technology exists to store, retrieve and manipulate information or data, whereas Operational Technology is mainly used to detect and cause change in operational processes.

Throughout the years, Information-based systems have seen significant technological advances in terms of speed and energy efficiency, as observed by Moore [181]. Typically an enterprise network will be comprised of systems such as personal workstations, various servers and peripheral devices such as printers. Due to their tasks of processing, transferring and modifying data, these require hardware that can perform tasks at efficient speeds.

	Information Technology	**Operational Technology**
Design Philosophy	Data Processing	Operational Control
Physical Design	Temperature-Oriented	Environment-Oriented
Power Requirement	Low-to-High	Low
Computation Power	Flexible	Limited to Intended Tasks
Uptime Requirement	Low-to-High	High

Table 2.3 Hardware differences between IT and OT

Memory-based hardware components such as storage drives, RAM, CPUs, and GPUs are, therefore, often heavily invested in to allow for larger resource loads. This flexibility allows IT hardware to withstand resource-intensive actions such as aggressive port scanning or vulnerability scanning for penetration tests for example.

Many devices are categorised under OT based on their specific role within the industrial process. Their function, which is to view, monitor, and control physical processes, directly impacts their hardware design and implementation. PLCs, for example, are devices that are designed to operate reliably within harsh environments (high temperature, wet conditions) for extended time-frames (several decades). As they are often designed with specific tasks in mind, their hardware architecture reflects this; being composed of a processor unit, power supply, an I/O interface, a communication interface and dynamic memory only, in most cases. Because of this, their hardware capability is often designed with durability and reliability in mind, meaning that components such as their processor units are often designed to operate at minimal power without interruption. The CPU resources, therefore, reflect this; for example, the SIEMENS S7-1200 PLC specifications indicate that its power consumption is 1.2A at 24V [251]. Consequently, the S7-1200's CPU processing time ranges between 0.085μs and 2.5μs per instruction depending on the operation type (bit, word or floating-point arithmetic), and its total available memory is 50kbyte and 1Mbyte for work and load memory, respectively. As we can see, the specifications for OT equipment differ significantly from the terms used for typical IT products, primarily because these are targeted towards automation engineers. Because of this, however, difficulties arise when attempting to directly compare these two system types directly, further demonstrating the IT/OT gap.

To summarise, due to the differences in functionality between IT and OT, the hardware specifications between these two system types also differ significantly, illustrated in Table 2.3. OT is designed to withstand harsh environmental factors for an extended time, whereas IT processes information as efficiently as possible. As such, OT is often considered more fragile, from a digital point of view, than IT.

2.2.3 Software/Firmware Characteristics

Similarly to the differences in hardware between OT and IT systems, there are many differences pertaining to software. This is also due to the underlying purpose of these devices. As IT systems are designed around data storage and manipulation, this is reflected in the software they run, which is both efficient and straightforward enough for mainstream use. Many IT devices such as personal computers or servers run commercially available or open-source Operating Systems (OS) such as Microsoft Windows, Apple's macOS, and GNU/Linux distributions, to name a few. These OSs offer an easy to understand Graphical User Interface (GUI) for everyday work projects or even for personal use. Specialised versions of these have also been developed with OSs designed specifically for servers, for example, with Microsoft Windows Server [168]. Over the years, these OSs have seen heavy investment in security with regular security patches and the development of built-in tools including proprietary anti-viruses or even a shift towards using biometric access control such as facial recognition or fingerprint readers instead of passwords [167]. Due to the popularity of these systems, new vulnerabilities and exploits are discovered regularly, resulting in the constant threat that machines are not updated regularly enough to keep up with newly discovered exploits despite vendors such as Microsoft being able to provide timely patches for newly discovered vulnerabilities. This was the case, for example, with the Wannacry attack on the UK National Health Service (NHS) in May 2017, where 80 out of 236 hospital trusts were affected due to having not made the appropriate updates to their Operating Systems, which was issued in March 2017 and advised by NHS Digital's CareCERT bulletin in April 2017 [256].

As opposed to IT systems, OT systems are designed to be used by specialised groups such as automation engineers. As such, industrial software also reflects this. PLCs, for example, are commonly programmed using Ladder Logic [22]. This specialised programming language represents written programs through graphical diagrams based on logic circuitry. Since its conception, other similar languages have been designed and standardised in IEC 61131-3 [116]. Ladder Logic provides a simple yet optimal method for controlling PLCs as, when programmed correctly, it is improbable to cause software or hardware crashes, which is critical when running inside of environments that require near-total uptime such as water supply stations or power plants [217]. Due to the critical aspect of ICS, engineers require quick and easy access to these if operational actions need to be modified or halted. Because of this, devices such as PLCs often run as root by default [218]. This has cause for several security concerns, especially considering that such PLCs are used as part of Critical Infrastructure networks. If an attacker were to gain access to these PLCs, they would have direct access to all of the PLCs' functionalities, resulting in severe operational impact.

	Information Technology	**Operational Technology**
Operating System	Open-Source, Commercially Available	RTOS, Ladder-Logic, Function Block Diagrams
Target Audience	Everyday Users	Specialised Engineers
Security Design	Security by-design	Operational Function by-design
Patching/Updates	Regular and Easy	Rare and Difficult

Table 2.4 Software differences between IT and OT

There has also been a push towards using Real-Time Operating Systems (RTOS) for various embedded systems within OT networks in recent years. These specialised OSs are used to serve real-time applications and process data without much or any buffering delay. This allows for the smooth operation of time-sensitive processes, which is often required within industrial networks [97]. However, a study on the security of RTOSs conducted by Yu et al. concluded that despite the advantages gained by using these, security against cyber attacks is still a significant concern in systems that make use of these specialised OSs [285].

Due to the criticality of requiring high uptime within OT environments, updating software for these can be challenging. A study presented at Black Hat USA found that the average time between disclosure of a vulnerability and detection of that same vulnerability within an OT network was 331 days at the time of the study [147]. Allowing attackers such an extensive time frame to develop exploits increases the risk associated with unpatched vulnerabilities and increases the potential impact that could be caused. Additionally, while IT systems can allow for timely security patches, not all OT systems can do the same depending on their functions. Power plants, for example, could schedule patching during the Summer, when not all systems are required for operations. However, providing software updates during the Winter could prove more challenging as the demand for electricity during that time is considerably higher [85, 264].

To summarise, the software differences between IT and OT, implemented based on their respective function, leads to a significant disparity in their security capabilities, as illustrated in Table 2.4. Furthermore, due to the only recent implementation of security features in OT software, critical features are often overlooked, leading to significant exposure to external threats.

2.2.4 Network Architecture and Protocol Characteristics

A common theme present throughout the comparison is that the difference in purpose between IT and OT systems undoubtedly leads to a difference in the characteristics of these systems. This also holds true for OT or IT network architecture and protocols. As IT networks were very early in their adoption of Internet Protocols, this naturally led to a need to defend

these systems against remote attacks [2]. Over the years, techniques and models have been developed to ensure that IT networks are highly resilient against both external and internal attacks. A significant number of network security reference architectures have since been developed and widely adopted, such as the Fortinet Network Security Reference Architecture [233] or through guidance provided by IBM [261]. The main objective of these architectures is to allow for good network flexibility while being resilient to threats. While there exist network reference architectures for OT environments such as the Purdue Enterprise Reference Architecture [40], these focus more on segmentation of network zones based on hierarchical function. Figure 1.1 illustrates the zones described in the Purdue Model. While there is little detail on how each layer can be secured at an individual level, the model uses a Demilitarized Zone (DMZ) to address security risks between the IT and OT zones. One such recommendation detailed by Cisco is to design the DMZ so that no traffic traverses directly through it, meaning that traffic must either terminate or originate in the DMZ. Historically, ICS/OT security also relied on the use of "air gapping" networks, making it so that external threats could not attack networks remotely. While this may have been a viable security policy a couple of decades ago, the introduction and implementation of smart networks, through IoT, makes it very difficult to truly have an air gap anymore [34].

It is, however, also worth noting the security capabilities of individual protocols used within OT networks. As most of these protocols were created at a time when security was not a primary concern, most of these have inherent vulnerabilities that can be exploited [89]. Attack vectors such as ARP spoofing, DNS cache poisoning or generally poor encryption need to therefore either be mitigated internally through firewall ruling or by implementing the secure version of the protocols under consideration, by using SSL/TLS for example [31]. This shift towards using more secure protocols is, however, not yet widespread within OT networks [20]. Despite some of the standardised protocols being updated to use TLS, such as Modbus TLS or OPC UA, a significant number of these protocols still do not make use of proper encryption or authentication such as PROFINET or CAN [280]. Modbus, as an example, is a protocol created by Modicon (now Schneider Electric) in the late 1970s and is a widely adopted OT protocol, mainly because of its simplicity and robustness. The Protocol Data Unit of the Modbus packet frame contains a function code and data payload. Packets can be handcrafted within any programming language to perform specific actions that could benefit an adversary, such as reading the device identification or even reading/writing directly to coils or registers. Table 2.5 provides several examples of Modbus Functions and their respective codes that could be used to perform malicious actions. This is mainly because, as an older protocol, Modbus lacks modern security features that would prevent attacks such as unauthenticated commands or replay attacks. This means that an adversary would be

Function Code	Function Description
1	Read Coils
4	Read Input Registers
5	Write Single Coil
6	Write Single Holding Register
14	Read Device Identification
15	Write Multiple Coils
16	Write Multiple Holding Registers
17	Report Slave ID

Table 2.5 Example Modbus Functions

able to control an operational process, such as through a traffic light system while making it appear from a connected HMI that the system under consideration was never altered. Similar security weaknesses exist within most ICS protocols, standard and proprietary alike. DNP3, for example, supports the use of a Direct Operation Function, which means that target devices can be actuated directly by the output points specified in the object of the received packet [1]. This can be done by anyone with access to the network or that can remotely communicate with a device using DNP3 due to the lack of authorisation control.

The documentation of these industrial protocols can be jarring at best, meaning that operators may be reluctant to upgrade implemented protocols to their more secure versions. The complete documentation available for the Common Industrial Protocol, for example, is over 1500 pages long [209]. Another challenge that industrial networks face is that many of these still make use of proprietary protocols such as S7COMM (Siemens) or Melsec/TCP (Mitsubishi Electrics) [172, 179]. While some have been updated to provide encryption and authentication, these protocols rely on their vendors for updates meaning that operators can be left with substantial delays between security updates.

In general, three mitigation techniques can be implemented to reduce the inherent risk of using OT protocols [20]. While it can be challenging to do so within a production environment, keeping firmware up to date can mitigate known network-based attack vectors such as arbitrary code execution or replay attacks. This, however, does not reduce the risk originating from undiscovered vulnerabilities (otherwise known as zero-days) but can help minimise their discovery. Most importantly, the implementation of proper network topology and segmentation, such as through using the Purdue Reference Architecture Model, can help reduce unauthorised access to industrial systems through means of layered defence [40]. Finally, incorporating Intrusion Detection Systems should also be done. However, as opposed

to IT, Intrusion Prevention Systems are discouraged for OT networks as their use could halt critical traffic within the process network due to false positives.

To summarise, due to the only recent requirements of implementing proper security control within OT networks, many gaps in security capabilities can be identified, as illustrated in Table 2.6. Despite a push to use more secure protocols, this is not yet comprehensive in practice, leaving industrial networks exposed to simple attack vectors due to the lack of security within communication protocols.

	Information Technology	**Operational Technology**
Network Architecture	Secure-by-Design	Safe-by-Design
Protocol Priority	Confidentiality	Availability
Protocol Security	Inherent	Optional
Authentication	Required	Little-to-None
Detection Systems	IDS/IPS	IDS

Table 2.6 Network and Protocol differences between IT and OT systems

2.2.5 Socio-Technical Characteristics

While the previous sections discussed technical differences between IT and OT, socio-technical aspects also play an essential part in the overall strategy for effectively securing assets against malicious actors. However, how these are implemented can differ between IT and OT.

One of the most well-known concepts for implementing cyber security controls and policies, for example, is the Confidentiality-Integrity-Availability (CIA) triad. The primary goal of organisations, in most cases, is to maintain the confidentiality, integrity and availability of information, ensuring complete coverage of cyber security capabilities. The ordering of this model reflects the prioritisation of these attributes, primarily for IT. As such, the preservation of confidentiality is overall allocated more priority over integrity and availability. This means that, in the event of a cyber attack, IT organisations will allocate more resources to ensuring the preservation of confidentiality than availability. This, however, is not the case for OT environments. Due to their time-critical nature, any change in availability can have detrimental consequences. For example, if the availability of a Safety PLC (part of a Safety Instrumented System) were to be compromised, this could most likely lead to an overall loss of safety. For this reason, OT security controls and policies prioritise the conservation of availability (and consequently integrity), whereas IT security controls and policies prioritise the conservation of confidentiality.

An observation noted during several surveys on ICS/OT cyber security [276, 281, 77] was the apparent lack of cross-disciplinary skills and communication between OT and cyber security. OT engineers, historically, were not required to be knowledgeable in security. However, with the rapid technological changes that have been introduced to this domain, this is now no longer the case. Because of this, senior engineers have been required to review their own practices from the ground up, leading to a significant disparity in skill and a shake-up in standard practices. To this day, an OT engineer is likely to firstly view a cyber incident as a technical fault before considering the event as a cyber attack, leading to a delay in appropriate response and recovery actions [264].

The use of standards and guidelines was also found to be more mature for IT than it is for OT. While a plethora of guidance is available for aiding OT operators in assessing and improving their cyber security capabilities, the disparity of topics discussed within these could leave operators with a less than complete picture, resulting in a potential gap in the implementation of security controls and policies [264]. To this end, existing standards and guidelines often lack the required tooling and frameworks for proper implementation within OT environments, instead applying information-based strategies as opposed to function.

This discrepancy can also be observed for certifications related to cyber security, especially security testing. At present, no accreditations exist to certify that an individual meets a specific level of understanding and expertise when conducting security engagements within industrial environments. A plethora of these exist for traditional IT, such as the Offensive Security Certified Professional, the Certified Ethical Hacker or the GIAC Penetration Tester certifications, but little currently exists for ICS/OT Penetration Testing [210, 68, 86]. The SANS Institute and ISA used to both offer SCADA and ICS penetration testing courses, but these are short in length, are not recognised by governing bodies, and are currently unavailable for enrolment [235, 124]. Instead, high-level courses are listed which provide a general overview of OT cyber security [234, 123]. Because of these challenges, only a small set of individuals are officially qualified to provide adversary-centric security tests for operators of OT, especially CNI.

To summarise, socio-technical aspects of security for OT has been observed to be significantly less mature than within IT, as illustrated in Table 2.7. Because of this, a significant knowledge gap exists between essential cyber security concepts and OT engineering, leading to a disparity in cyber security capabilities between IT and OT.

2.2.6 Summary

From an asset management point of view, the differences between IT and OT can be grouped into four distinct categories: hardware; software; network; and socio-technical differences.

	Information Technology	Operational Technology
CIA Triad Prioritisation	Confidentiality	Availability and Safety
Security Culture	Mature	In-Progress
Use of Standards and Guidelines	Comprehensive	Limited
Certifications	Widely Available	Little-to-None

Table 2.7 Socio-Technical differences between IT and OT

Summarised in Table 2.3, the design of both IT and OT hardware is directly correlated to their function within their environment. Because the goal of using IT is to store, process and exchange information, its hardware is designed to enable this as efficiently as possible, being flexible in processing tasks. In contrast, OT is designed for viewing, monitoring, and controlling operational processes, leading to their hardware design focusing on environmental resilience, high up-time, and cost-efficiency, often limited to processing highly specialised tasks. Similarly, the software of these system types is designed in consideration of their end-users, summarised in Table 2.4. Because IT is commonly used throughout different domains, its software is flexible and designed to be simple to understand. On the other hand, OT software is explicitly designed to be used by automation engineers, making skill transference difficult between the two domains. Additionally, a focus on safety is prioritised over security due to the critical nature of OT environments; this makes security updates challenging to implement and, therefore, uncommon due to the high up-time requirements of these systems. The critical nature of OT environments also affects the approach to their network architecture and the design of industrial protocols, summarised in Table 2.6. Because of the time and safety-critical aspects of OT networks, protocols used within them often lack fundamental security implementations, such as access control and encryption, despite being widely adopted in IT systems. These technical differences directly influence the sociological factors behind the implementation of security controls and the culture within IT and OT environments, further demonstrating the gap between IT and OT, as summarised in Table 2.7.

2.3 Conclusion

In conclusion, it is clear that ICS/OT have been increasingly targeted by cyber attacks in recent years. These attacks can have severe consequences, ranging from financial losses and reputational damage to physical harm and even loss of life.

The main difference between IT and OT lies in their purpose and scope. While IT is primarily concerned with the management and processing of data, OT is focused on the

control and monitoring of physical processes and devices. This difference has significant implications for cyber security, as OT systems are often more complex and less secure than IT systems. To effectively protect ICS/OT systems from cyber attacks, it is essential to understand the unique characteristics and vulnerabilities of both IT and OT. This requires a holistic approach that considers the integration and alignment of both systems, as well as the implementation of appropriate cyber security measures.

Because of this dramatic increase in cyber attacks targeting CNI by skilled threat actors such as nation-state funded organisations, effectively preparing to respond to and recover from these attacks has become essential; especially considering that modern adversaries are willing to invest a significant amount of resources to achieve their goal, regardless of the cost [58]. Chapter 3 therefore aims to present existing practices for cyber incident response and recovery and identify gaps in both theory, through analysis of standards and guidelines, and practice, through stakeholder interviews. From there, clear requirements can be defined to determine areas of improvement in the context of incident response and recovery.

Chapter 3

Analysis of Current Industry Response and Recovery Practices

As discussed in Chapter 2, an increasing series of attack targeting ICS/OT have been observed, especially within the last decade. These attacks have acted as a catalyst for change in how we consider cyber defence within an industrial context. With organizational drivers and ever-evolving technical capabilities pushing the boundaries of safety and security to meet end-user goals, adoption, and increasing maturing of cyber security as a whole is becoming essential.

The importance of cyber security in CNI has not gone unnoticed by governments on an international level, many of whom have introduced strategies to drive change. For example, in 2016, EU member states introduced the NIS Directive [72]. In the UK, this was followed in 2016 by the creation of a NCSC [186], whose core role is to provide cyber security advice and support for public and private sector organizations, including focused advice on NIS Directive compliance for CNI [185]. Similarly, in 2013 the United States of America assigned the NIST the task of providing guidance on cyber security for CNI [213] Guidance from organisations such as the NCSC and NIST, primarily focus on five key principles; identify, protect, detect, respond, and recover [187]. The latter two of which (respond and recover) can be considered as the last line of defence, designed to limit the impact of a cyber incident and promote a prompt recovery.

Although cyber incident response and recovery is crucial in most cyber security strategies, it is less explored than other areas. Given its last line of defence status, it presents a critical component that must be well understood by CNI operators. This chapter provides an analysis of existing ICS/OT focused standards and guidelines to identify the construct of response and recovery processes, their level of coverage, and potential challenges faced when using these documents. This analysis acts as a foundation in a set of semi-structured interviews with CNI

operators and regulators to better understand current response and recovery practices, the use of existing standards and guidelines, and any associated challenges. These two studies then form a set of requirements from which a framework, provided in Chapter 4 has been designed to support CNI operators in developing response and recovery capabilities through better use of standards and guidelines.

3.1 Related Work

Over the last decade, there has been an increasing volume of research activity targeted towards the holistic improvement of cyber security deficiencies within an ICS/OT context. However, the field of cyber incident response and recovery has seen less focus [139], compared with risk assessment, for example.

Several works [98, 56, 289, 255] have explored historic cyber attacks against ICS/OT. The work of Hassanzadeh et al., [98], for example, includes coverage of response, remediation, and lessons learnt. This can be used to better understand adversaries, their actions, and the actions of targeted organisations. All of which can support operators in the development of their own cyber security capabilities.

There exists a broad range of work on intrusion detection for ICS/OT, all of which contribute towards the information set available to operators during initial incident response activities. For example, Jardin et al., [133] propose a non-invasive active monitoring approach, in contrast to more traditional passive techniques [90]. Taking an alternative approach, Urbina et al., [277] explore physics-based attack detection as a mechanism by which the impact of stealthy attacks can be minimised. A new metric is introduced to measure the impact of stealthy attacks, followed by a proposed combination and configuration of detection schemes towards stealthy attack mitigation. While Casalicchio and Gualandi [33] focus on the detection of changes to control logic. This is described as a self-protecting architecture for cyber-physical systems.

Going one step further from baseline detection techniques, Piedrahita et al., [222] apply software defined networking into their detection and response system. This allows for automated network reconfiguration as a means of mitigation during an incident. With similar motivations, Ullah et al., [275] model an intrusion response system through the consideration of diverse attacker strategies. This model explores potential attack pathways, from which a response mechanism is designed to restrict attacker opportunities. Similarly, Cook et al., [43] define a seven-stage triage process to determine areas of priority where an attack's impact would be most significant.

Practical applied recovery work is also well covered. The work of Khalili et al., [143] for example, presents a recovery scheme for ICS/OT, focusing on reducing the Mean Time To Recovery (MTTR). This work describes the use of physical backup hardware in recovering a system to its pre-attack state. Sesaki et al., [236] also explore system recovery and propose a novel approach by using a fallback and recovery ICS/OT, the Fallback Control System (FCS). The FCS is not networked, and isolates controlled objected away from networked devices to manage them safely in isolation.

Butts and Glover [26] explore and describe limitations in current ICS/OT security training. As part of this discussion, they describe the need for training to carry out response activities, develop training facilities with real-world environments, multiple interconnected systems, etc. An example of a response coordination syllabus is also outlined. Hirai et al., [103] begin to address these challenges by exploring incident response roles and responsibilities and introducing a framework for cyber incident response training. Further, Antonioli et al., [10] explore gamifying security training.

Cyber exercising is seen as a form of training, with the work of Asai et al., [15] proposing a framework that discusses exercise design, evaluation, and management. A practical exercise is provided as a means of validating the proposed framework. This is a theme explored by others in the creation of exercise platforms/testbeds for a variety of related activities [144, 227, 8].

ICS/OT forensics has also been explored, including forensic readiness spanning data sources and tooling [69, 3, 19], case studies [279], and overarching forensic architectures [284]. Line at al., [158] engage with industry stakeholders to explore cyber situation awareness. This work focuses on comprehension of the current situation and understanding impact, situation evolution, attacker behaviour, and cause. From this, the authors provide a set of five recommendations (exercise, prepare for social engineering attacks, physical network separation, deploy anomaly detection, and use regulation as a means of ensuring improvements) focused on detection and response.

The work of He et al., [99] propose an ICS/OT incident response decision framework across three phases (Descriptive, Predictive, and prescriptive). Jaatun et al., [132] propose a framework for incident response management in the petroleum industry. This work also includes engagement with industry stakeholders across multiple studies on incident response, risk and vulnerability assessment, security challenges at an installation, overall project findings, etc. These provide motivation and input into the resulting framework. The framework provides a high-level overview of factors one should consider as part of their overall response and recovery capabilities. However, it is a combination of just two (now outdated) standards and guidelines “with increased emphasis on proactive preparation and

reactive learning". Line et al., [159] also interview industry stakeholders within an industrial context to better understand security incident management. This is focused on comparing small to large organisations but offers valuable insight into key challenges.

To summarise, cyber incident response and recovery has received limited attention from a holistic perspective, a critical gap noted by others [139]. While, collectively, existing literature spanning intrusion detection, historic attack analysis, training, exercising, etc., all contribute towards improvements in cyber incident response and recovery capabilities, a higher-level understanding of core requirements is still required. Furthermore, additional engagement with industry stakeholders and further understanding of challenges faced when using official standards and guidelines could add further towards understanding and addressing key challenges.

Industry standards and guidelines currently offer the most holistic and well-established view on response and recovery requirements. However, no comprehensive analysis of these has been undertaken to explore commonality in approach, coverage of key themes, use of technical vs non-technical content, etc. The following section therefore provides an analysis of these standards and guidelines to better understand each fundamental response and recovery phase and associated sub-phases; as well as to identify gaps in these documents that could negatively affect in-practice response and recovery capabilities.

3.2 Analysis of ICS/OT Cyber Incident Response and Recovery Standards and Guidelines

3.2.1 Document Selection

The following subsections provide an analysis of selected government and industry standards and guidelines to support the development and delivery of cyber incident response and recovery capability used for this analysis. This guidance is primarily targeted at those responsible for the continued safe operation of ICS/OT. Initial exploration focuses on UK-centric guidance and any supplementary documentation (i.e. referenced materials). International guidance is then investigated with a focus on North America and France, selected based on their accessibility (i.e. Open to the public and written in English), and their global nuclear energy presence, acting as an indicator of required cyber security guidance for one of the most critical elements of CNI [122]. The objective of this analysis is to explore the contents of standards and guidelines that discuss ICS/OT incident response and recovery and identify the challenges that operators of ICS/OT face when consulting different publications.

Guidance/Standard	Organisation	References
NCSC Cyber Assessment Framework	NCSC	[194]
DWI Cyber Assesment Framework	DWI	[67, 66]
10 Steps: Incident Management	NCSC	[192]
Security Assessment Principles (SyAPs)	ONR	[212]
Preparation for and Response to Cyber Security Events Technical Assessment Guide	ONR	[214]
HMG Security Policy Framework	HMG	[104]
Operational Guidance 86	HSE	[104]

Table 3.1 Overview of Selected UK Guidance and Standards

Guidance/Standard	Organisation	References
Nuclear Security Fundamentals	IAEA	[110]
Nuclear Security Series 17	IAEA	[109]
Nuclear Security Series 23-G	IAEA	[111]
Good Practice Guide for Incident Management	ENISA	[111]
Computer Security Incident Response Team FAQ	Carnegie Mellon University	[32]
Incident Handler's Handbook	SANS	[151]
Security Consensus Operational Readiness Evaluation	SANS	[267]
CIS Critical Security Controls	CIS	[38]
SP 800-61	NIST	[36]
SP 800-53	NIST	[202]
Cyber Security Incident Response Guide	CREST	[47]
ISO/IEC 27001/27002	ISO/IEC	[129, 130]
ISO/IEC 27035:2016	ISO/IEC	[127, 128]
IEC 62443 Series	IEC	[115, 119]

Table 3.2 Overview of Selected Supplementary Guidance and Standards

A high level summary of the selected guidance documents can be found in Tables 3.1, 3.2, and 3.3 (UK guidance, supplementary/reference guidance and international guidance respectively). For more detail on the specific contents of each resource, a detailed summary of each document is provided in the following subsections.

3.2.2 Overview of UK Guidance

NCSC Cyber Assessment Framework (CAF)

Created in response to the NIS Directive, the CAF consists of four objectives, each focusing on a different stage of an organisation's security planning [194]. Objective D relates to guidance on response and recovery and is broken down into two sub-objectives: Response and Recovery Planning and Lessons learnt. The CAF also recommends consulting additional external resources [36, 127, 47].

Guidance/Standard	Organisation	Country	References
ISO/IEC 27019:2017	ISO/IEC	N/A	[131]
RG 5.71	NRC	USA	[206]
NEI 08.09	NEI	USA	[207]
Framework for Improving Critical Infrastructure	NIST	USA	[187]
SP 800-82	NIST	USA	[266]
SP 800-83	NIST	USA	[260]
SP 800-100	NIST	USA	[23]
CIP-008-06	NERC	USA	[197]
REGDOC-2.5.2	CNSC	Canada	[41]
Managing Cyber Security for Industrial Control Systems	ANSSI	France	[9]

Table 3.3 Overview of Selected International Guidance and Standards

Drinking Water Inspectorate (DWI) Cyber Assessment Framework

The DWI have published their own guidance tailored towards the water sector in the UK. Based on the NCSC's CAF, it aims to provide operators with a framework for managing cyber security risks and incidents that could impact drinking water quality or availability. Furthermore, it allows the DWI to assess operators' security measures for compliance with the NIS Directive [67]. The DWI CAF is constructed around four top-level objectives, objective D being related to response and recovery [66]. This guidance recommends consulting additional external resources [129, 115, 202].

NCSC 10 Steps: Incident Management

The NCSC 10 Steps for incident management provides a light-weight resource covering key considerations aligned to incident response and recovery activities [192]. These include establishing a response capability, providing training, and usage of lessons learnt.

Office for Nuclear Regulation (ONR) Security Assessment Principles (SyAPs)

SyAPs aid regulatory judgements and recommendations when undertaking assessments (for compliance) of nuclear facilities [212]. The assessment principles contain ten Fundamental Security Principles (FSyPs), two of which are directly relevant to cyber incident response and recovery (FSySP 7 and 10). These cover the following topics: Counter-Terrorism Measures, Emergency Preparedness, Response Planning, Testing and Exercising of the Security Response, and Clarity of Command, Control and Communications Arrangements During a Post Nuclear Security Event.

ONR Preparation for and Response to Cyber Security Events Technical Assessment Guide (TAG)

This TAG provides guidance for ONR inspectors' use covering eleven topics related to cyber security event response [214]. While TAGs explicitly state that they are not a resource for demonstrating adherence to SyAPs, they can provide additional insight into what operators' high-level goals should be. This guide also recommends consulting external resources [110, 109, 111, 71, 32, 151, 267, 38].

His Majesty's Government (HMG) Security Policy Framework

The HMG Security Policy Framework covers several topic areas, from culture and awareness to risk management and personnel security [104]. Although brief, one section describes requirements when preparing for, and responding to, security events. This is discussed using generic, non-cyber terminology.

Health and Safety Executive (HSE) Operational Guidance (OG) 86

OG 86 is closely aligned to the NCSC CAF and is formed around its core security objectives and corresponding principles [106]. Discussion on cyber incident response and recovery is present throughout this guide. Guidance surrounding cyber incident response and recovery is provided in direct alignment to CAF objective D. This can be summarised as the development of a clear and concise, well-articulated cyber incident response plan. OG 86 also recommends consulting additional external resources [115, 129].

3.2.3 Overview of Supplementary Guidance

International Atomic Energy Agency (IAEA) Nuclear Security Fundamentals

The IAEA Nuclear Security Fundamentals outlines 12 essential elements required to support a state's nuclear security regime [110]. Cyber security is only mentioned once within this document, linked to a requirement on assurance activities. Essential element 11 relates directly to response (i.e. planning for, preparedness for, and response to, a nuclear security event).

IAEA Nuclear Security Series (NSS) 17

NSS 17 is designed to guide operators in establishing and improving programmes of work to protect computer systems, networks, and other (critical) digital systems responsible

for the safe and secure operation of nuclear facilities [109]. Specific details on cyber incident response and recovery are limited to generic guidance, such as describing relevant responsibilities and response planning.

IAEA NSS 23-G

The objectives of NSS 23-G [111] are defined over four areas: establishing a framework for ensuring the confidentiality, integrity, and availability of sensitive information; identifying sensitive information; considerations for sharing/disclosing sensitive information; and guidelines/methodologies. Therefore, its ties to cyber incident response and recovery are limited; however, content such as that found in Annex 2 (i.e. examples of sensitive information) could be used when categorising information related to "contingency and response plans and exercises".

ENISA Good Practice Guide for Incident Management

While not directed towards ICS/OT, this guide provides a comprehensive discussion on cyber incident management for conventional IT systems [71]. Covered topics include response and recovery by explaining the incident handling process and basic codes of practice.

Carnegie Mellon University - Computer Security Incident Response Team FAQ

This FAQ provides a high-level discussion on CSIRTs. Although not targeted towards ICS/OT, it acts as a helpful reference point in understanding core CSIRT requirements [32].

SANS Incident Handler's Handbook

The SANS Incident Handler's Handbook details key phases of incident response and recovery, their purpose, tools that can be used to support them, etc. [151]. While this is not ICS/OT specific, it provides a comprehensive discussion on response and recovery broken down into the following core sections: Preparation, Identification, Containment, Eradication, Recovery and Lessons Learnt.

SANS Security Consensus Operational Readiness Evaluation (SCORE)

The SANS SCORE security checklist is highly summarised in the form of six bullet points, each corresponding to the six steps presented within the Incident Handler's Handbook. It is designed to support all forms of incidents, including those from Advanced Persistent Threats (APT) [267].

The Center for Internet Security (CIS) Critical Security Controls (CSC)

CIS CSC presents 20 security controls [38]. Although defined as controls, these are more closely linked with high-level groups/objectives, to which mapping against the NIST Cyber Security Framework is performed. CSC 10, 19 and 20 discuss response and recovery topics covering guidance for both small and large organisations.

NIST Computer Security Incident Handling Guide (SP 800-61)

SP 800-61 details the need for incident prioritisation, stating that the handling and subsequent recovery of systems affected by these incidents should be determined by the potential impact on service functionality and information integrity [36]. A focus is placed on exploring methods for ensuring essential service continuity and impact mitigation.

NIST SP 800-53

SP 800-53 provides a "catalogue of security and privacy controls for federal information systems and organisations to protect organisational operations and assets, individuals, other organisations, and the Nation from a diverse set of threats including hostile attacks, natural disasters, structural failures, human errors, and privacy risks" [202]. This includes twenty mandatory controls, mapped to ISO/IEC 27001 [129], for securing assets, including response and recovery.

CREST Cyber Security Incident Response Guide

This guide is split into three core areas: preparing for, responding to, and recovering from a cyber security incident [47]. Each of these areas contains a step by step guide offering potential avenues for an organisation to follow during incident response, including methods for identifying potential incidents, conducting triage, and effectively containing and recovering from a state of containment.

BS EN ISO/IEC 27001/27002

ISO/IEC 27001 provides non-technical guidance for implementing and maintaining systems that are well protected from cyber threats, including a table in Annex A listing all the objectives that an asset owner should achieve [129]. Section A.16 of this table refers to incident management and consists of the following objectives: Responsibilities and Procedures, Incident Reporting, Vulnerability/Weakness Reporting, Event Assessment, Incident Response, Lessons Learnt, Evidence Collection. These objectives are described in more

detail within ISO/IEC 27002, which serves as a "best practices" guidance for implementing the requirements in ISO/IEC 27001 [130].

BS EN ISO/IEC 27035:2016

ISO/IEC 27035 serves as a reference for fundamental principles designed to ensure that the correct tools, techniques and methods are appropriately selected in the event of a cyber incident. Part 1, Principles of incident management, presents fundamental concepts of information security incident management. These concepts are combined with principles from the five phases of response and recovery: detecting, reporting, assessing and responding to incidents, and applying lessons learnt [127]. Part 2, Guidelines to plan and prepare for incident response, describes how to plan and prepare for cyber incident response and recovery. This covers the "Plan and Prepare" and the "Lessons Learnt" phases presented in Part 1 of the standard [128].

BS EN IEC 62443 Series

The IEC 62443 catalogue defines procedures for implementing secure ICS/OT (referred to as Industrial Automation and Control Systems). However, while the entirety of the catalogue was recommended by UK guidance, due to paywall restrictions, only parts 2-1 and 4-2 of the series were selected. Part 2-1 of this series provides guidance for establishing an ICS/OT security program, including planning for incident response and recovery [115]. Part 4-2 of the series describes the technical security requirements for ICS/OT components, including guidance on how to ensure that systems respond promptly to security violations by alerting the appropriate personnel and reporting details on the violation [119].

3.2.4 Overview of International Guidance

BS EN ISO/IEC 27019:2017

ISO/IEC 27019 provides guidance to fulfil the objectives set out in ISO/IEC 27001 and 27002 for ICS/OT within the energy utility industry [131]. This is similar to that provided in ISO/IEC 27001, with subtle modifications to better suit ICS/OT.

Nuclear Regulatory Commission (NRC) RG 5.71 (USA)

RG 5.71 provides a comprehensive overview of cyber incident response and recovery guidance for nuclear operators [206]. Guidance is provided under high-level requirements for establishing a cyber security plan concerning incident response and recovery.

Nuclear Energy Institute (NEI) 08.09 (USA)

NEI 08.09 is closely linked to NRC RG 5.71 [207]. Response and recovery activities/requirements are discussed across multiple high-level topic areas surrounding contingency planning.

NIST Framework for Improving Critical Infrastructure 2018 (USA)

This framework focuses on improving cyber security risk management for CNI [187]. It provides a standard organisational structure for multiple cybersecurity approaches by assembling standards, guidelines, and practices into one document. Five core functions are defined, two of which are related to response and recovery.

NIST SP 800-82 (USA)

SP 800-82 provides guidance on securing ICS/OT. It presents a general overview of system architectures, associated vulnerabilities, and recommendations on how to counteract these in order to reduce the associated risk [266]. ICS-specific response and recovery guidelines include Incident Detection, Incident Classification, Response Actions, and Recovery Actions.

NIST SP 800-83 (USA)

Based on SP 800-61, SP 800-83 provides a Guide to Malware Incident Prevention and Handling for Desktops and Laptops. Although not ICS/OT specific, it is intended to help operators understand and mitigate risks associated with malware incidents, including associated practical guidance on response activities [260].

NIST SP 800-100 (USA)

SP 800-100 provides high-level guidance for management personnel tied to general information security themes, including risk management, service acquisition, and planning [23]. Guidance on response and recovery includes topics such as Incident Preparation, Incident Prevention, Incident Eradication, Incident Recovery, and Post-Incident Activities.

North American Electric Reliability Corporation (NERC) CIP-008-06 (USA)

Targeted towards power systems in North America, CIP-008-06 encompasses cyber security incident reporting and response planning requirements and associated recommendations. Its purpose is to “mitigate the risk of reliable operation of the Bulk Electric System as the result of a cyber security incident by specifying incident response requirements” [197].

Canadian Nuclear Safety Commission (CNSC) REGDOC-2.5.2 (Canada)

Cyber security requirements feature throughout this document within discussions on design management, design documents, and the instrumentation of the control life-cycle. One section is dedicated to cyber security under "Robustness Against Malevolent Acts" [41]. While this section provides a set of guiding principles focused on ties to safety, the inclusion of cyber specific incident response and recovery guidance is limited.

ANSSI Managing Cyber Security for Industrial Control Systems (France)

Coverage of cyber incident response and recovery activities within this document is limited, appearing briefly as part of a discussion on defence-in-depth strategies and across one section focusing on the "Incident Handling Alert Chain". This section is brief, with best practices established in the form of three questions [9].

3.2.5 Critical Analysis of Standards and Guidelines

Across the aforementioned thirty-one standards and guidelines, a range of cyber incident response and recovery phases/sub-phases can be identified. While as a collective, the thirty-one standards and guidelines provide a comprehensive guidance base, should individual resources be used in isolation, a less than complete picture of requirements could be formed and, therefore, misused in practice. Consequently, The following subsections provide an analysis of key phases and sub-phases and their coverage across each independent resource.

Methodology

When assessing the effectiveness of current guidance for ICS/OT cyber security response and recovery, relevant requirements must first be identified. During initial read-through of the selected documents, response and recovery phases/sub-phases were identified and extracted for use as a base for the analysis. Table 3.4 takes each of these identified phase/sub-phase and aligns it to a criteria set. Additional criteria of technical and non-technical factors are also included, allowing for a clearer understanding of each resource's target audience. The resources from Tables 3.1, 3.2, and 3.3 were then independently compared against this criteria set, ensuring a structured analysis could be undertaken.

Results

The analysis results have been compiled into Tables 3.5 and 3.6, offering a high-level snapshot of criteria coverage within each resource. In addition, key findings can be broken down

Requirement Type/Phase	Requirement	Criteria
Type	Non-Technical (NT)	Information provided is non-technical.
	Technical (Tec)	Information provided is technical.
Planning	Roles and Responsibilities (RR)	Contains information on assigning/defining roles and responsibilities.
	Response Planning (RP)	Contains information on response plan documenting.
	Criticality Assessment (CA)	Contains information on identifying and assessing key assets and infrastructure in terms of criticality.
	Threat Analysis (TA)	Contains information on conducting a continuous threat analysis for remediating identified vulnerabilities and minimising attack vectors.
	Risk Management (RM)	Contains information on creating and consulting risk management documents.
Preparation	Training (Tra)	Contains information on personnel training - including response team training/awareness training.
	Regular Testing and Auditing (RTA)	Contains information on testing and auditing- this includes red team exercises, penetration tests, and automatic testing.
	Incident Detection (ID)	Contains information on incident detection mechanisms.
Mid-Incident	Resource Availability (RA)	Contains information on resource allocation and accessibility in the event of a cyber incident (physical and non-physical resources).
	Incident Reporting (IRep)	Contains information on reporting incidents to the appropriate personnel (internal/external).
	Incident Containment (IC)	Contains information on procedures that should be implemented for containing the damage caused by an incident.
	Incident Eradication (IE)	Contains information on procedures that should be implemented for eradicating incidents.
	Incident Recovery (IRec)	Contains information on procedures that should be implemented for recovering from an incident.
	Evidence Collection/Handling (EC)	Contains information on evidence collection for use by external authorities.
	Public Relations Management (PRM)	Contains information on public information disclosure management.
Post-Incident	Root Cause Analysis (RCA)	Contains information on post-incident analysis; used to determine the root cause of the incident.
	Lessons Learnt (LL)	Contains information on lessons learnt from past incidents for improving current defensive capabilities.

Table 3.4 Requirements and Criteria for Document Analysis

across the four primary phases (Planning, Preparation, Mid-Incident, and Post-Incident) as follows:

- **Planning -** The majority of investigated guidance (∼80%) discusses the importance of response plan documenting and role and responsibility assignment. There is, however, minimal discussion on Criticality Assessment and Threat Assessment (∼53% and ∼63% respectively). Also, Risk Management is inconsistently discussed (∼53%).
- **Preparation -** The need for training is well covered (71%). However, discussion on testing and auditing, in addition to incident detection, is inconsistent (∼59% and ∼66%, respectively). This may be due to its inclusion within a larger series. For example, The NCSC CAF Objective D does not cover incident detection, as this is covered in Objective C [195].
- **Mid-Incident -** Incident Reporting, Containment, Eradication and Recovery are well covered (∼88%, ∼69%, ∼72%, and ∼81% respectively). However, Resource Availability guidance is limited (∼34%). This level of coverage is surprising, as most post-incident activities are highly dependent on resource availability. If resources (human and non-human) are incorrectly allocated or unavailable, this can adversely affect an incident's impact. Additionally, Evidence Collection/Handling and Public Relations Management coverage is limited (∼31%, and ∼19%, respectively). This also is a cause for concern, especially considering the importance of these activities. For serious incidents, the collection and preservation of evidence for authorities is essential. Any accidental tampering of evidence during response and recovery activities could seriously affect the corresponding investigation. Similarly, maintaining an honest and trustworthy reputation with the general public is crucial, as this can affect operations in the long term.
- **Post-Incident -** Although Lessons Learnt are well covered (∼66%), phases that directly impact the quality of this remain inconsistently discussed. Without consistent discussion of Root Cause Analysis (∼31%), limited guidance is, in reality, available for ensuring that the output from Lessons Learnt is thorough enough.

3.2.6 Summary

The majority of analysed resources contain high-level details, with only ∼54% providing technical guidance. Since ICS/OT implementations can differ between environments, hardware-specific technical guidance is not always recommended. However, due to the

	Type		Planning					Preparation		
Guidance/Standard	**NT**	**Tec**	**RR**	**RP**	**CA**	**TA**	**RM**	**Tra**	**RTA**	**ID**
SyAPs (ONR)	✓		✓	✓	✓	✓	✓	✓	✓	✓
TAG (ONR)	✓		✓	✓	✓		✓	✓	✓	✓
CAF - Objective D (NCSC)	✓		✓	✓	✓		✓	✓	✓	
10 Steps: Incident Management (NCSC)	✓		✓	✓			✓	✓	✓	
OG 86 (HSE)	✓	✓	✓		✓		✓			
Security Policy Framework (HMG)	✓				✓	✓	✓			✓
CAF (DWI)	✓		✓	✓	✓		✓	✓	✓	
Cyber-Security Incident Response Guide (CREST)	✓	✓		✓	✓	✓				✓
Good Practice Guide for Incident Management (ENISA)	✓		✓	✓			✓			
Nuclear Security Fundamentals (IAEA)	✓		✓	✓	✓			✓	✓	✓
NSS 17 (IAEA)	✓	✓	✓	✓	✓	✓	✓	✓		
NSS 23-G (IAEA)	✓	✓	✓	✓			✓	✓	✓	
IEC 62443 (Parts 2-1 and 4-2)	✓	✓	✓	✓				✓	✓	✓
ISO/IEC 27001/27002	✓		✓	✓		✓			✓	
ISO/IEC 27035	✓	✓		✓		✓	✓	✓		✓
ISO/IEC 27019	✓		✓	✓		✓			✓	
RG 5.71 (NRC)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
NEI 08.09	✓	✓	✓	✓		✓	✓	✓	✓	✓
NIST Framework	✓		✓	✓	✓	✓	✓	✓	✓	✓
NIST SP 800-53	✓	✓	✓	✓		✓	✓	✓	✓	✓
NIST SP 800-82		✓		✓		✓				✓
NIST SP 800-83	✓	✓		✓		✓		✓		✓
NIST SP 800-61	✓	✓		✓		✓		✓		✓
NIST SP 800-100	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
CIP-008-06 (NERC)	✓		✓	✓	✓				✓	✓
CSIRT FAQ (Carnegie Mellon University)	✓		✓	✓		✓		✓	✓	
Incident Handler's Handbook (SANS)	✓	✓	✓	✓	✓			✓		✓
SCORE (SANS)	✓	✓	✓	✓		✓				✓
Critical Security Controls (SANS)	✓		✓	✓		✓		✓	✓	✓
REGDOC-2.5.2 (CNSC)	✓	✓	✓	✓	✓	✓		✓	✓	✓
Managing Cyber Security for ICS (ANSSI)	✓	✓	✓	✓	✓	✓	✓	✓		✓

Table 3.5 Document Analysis Results (Part One)

	Mid-Incident							Post-Incident	
Guidance/Standard	**RA**	**IRep**	**IC**	**IE**	**IRec**	**ECH**	**ERM**	**RCA**	**LL**
SyAPs (ONR)	✓	✓							
TAG (ONR)		✓	✓	✓	✓	✓			
CAF - Objective D (NCSC)	✓	✓	✓	✓	✓			✓	✓
10 Steps: Incident Management (NCSC)	✓	✓			✓				✓
OG 86 (HSE)	✓	✓	✓	✓	✓			✓	✓
Security Policy Framework (HMG)		✓				✓			
CAF (DWI)	✓	✓	✓	✓	✓				✓
Cyber-Security Incident Response Guide (CREST)	✓	✓	✓	✓	✓			✓	✓
Good Practice Guide for Incident Management (ENISA)		✓							
Nuclear Security Fundamentals (IAEA)	✓								
NSS 17 (IAEA)					✓				
NSS 23-G (IAEA)		✓				✓	✓	✓	✓
IEC 62443 (Parts 2-1 and 4-2)		✓	✓	✓	✓				✓
ISO/IEC 27001/27002		✓	✓	✓	✓	✓			✓
ISO/IEC 27035		✓	✓	✓	✓				✓
ISO/IEC 27019		✓	✓	✓	✓	✓			✓
RG 5.71 (NRC)		✓	✓	✓	✓		✓	✓	✓
NEI 08.09		✓	✓	✓	✓				✓
NIST Framework	✓	✓	✓	✓	✓		✓	✓	✓
NIST SP 800-53		✓	✓	✓	✓		✓		
NIST SP 800-82			✓	✓	✓				
NIST SP 800-83		✓	✓	✓	✓	✓		✓	✓
NIST SP 800-61		✓	✓	✓	✓	✓		✓	✓
NIST SP 800-100)		✓	✓	✓	✓				✓
CIP-008-06 (NERC)		✓				✓		✓	✓
CSIRT FAQ (Carnegie Mellon University)	✓	✓	✓	✓	✓		✓		✓
Incident Handler's Handbook (SANS)	✓	✓	✓	✓	✓	✓			✓
SCORE (SANS)		✓	✓	✓	✓	✓	✓	✓	✓
Critical Security Controls (SANS)		✓	✓	✓	✓				
REGDOC-2.5.2 (CNSC)		✓	✓		✓				
Managing Cyber Security for ICS (ANSSI)					✓				

Table 3.6 Document Analysis Results (Part Two)

subject area's technical nature, a lack of technical guidance may present a challenge for operators during practical implementation.

Through the analysis of the selected resources, a lack of consistency has been highlighted. Although the core topics surrounding cyber security response and recovery activities are discussed in most, including Roles and Responsibility assignment and Response Plan Documenting, less-common topic areas, Evidence Collection or Public Relations Management, for example, appear irregularly. Concerns arise where operators are recommended to consult guidance that does not discuss these topics, leading to overlooked critical activities. While this analysis of standards and guidelines allows for operators and researchers alike to better understand the requirements needed to develop a comprehensive ICS/OT cyber incident response and recovery plan as well as identify which publications cover specific topics, the lack of consistency throughout available guidance highlights the need for amalgamation into a single resource, which operators can consult; ensuring complete coverage.

Having provided an analysis of a broad literature base to identify the current state of the art guidance for cyber incident response and recovery, the following sections detail the creation of synthetic cyber attack scenarios applied to a set of interviews with industry stakeholders. This provides a picture of current real-world practices and how the gaps identified in the resources discussed here can affect incident response and recovery capabilities in practice. The goal for this is to establish a base from both theory (standards and guidelines) and practice (stakeholder interviews) for use in the creation of a framework; detailed in Chapter 4.

3.3 Synthetic Scenario Development

To support engagement with industry stakeholders (see Section 3.4) and avoid findings being tied directly to real-world infrastructure, creating realistic synthetic attack scenarios is required. This presents a generalisable foundation on which all participants can openly discuss their approaches to cyber incident response and recover while ensuring neither the research team nor the participant cross sensitive information boundaries. Discussion of these scenarios also enables interview participants to discuss how specific guidance and guidelines, discussed in Section 3.2, could positively or negatively affect ICS/OT cyber incident response and recovery activities. The construction of these scenarios is outlined over the following subsections.

3.3.1 Historical Attacks

A set of historical attacks from Chapter 2 was reviewed to contextualise better the risk posed to ICS/OT and create realistic synthetic cyber attack scenarios for use in stakeholder engagement. These ranged from simplistic coincidental malware infections to more sophisticated targeted attacks. While there exist many a discussion across media outlets concerning cyber attacks targeting ICS/OT, many of these describe generic scanning tool traffic seen on the Internet. For this reason, the attacks listed here were more focused, with evidence of witnessed impact. Each contains a common name, year of occurrence, initial access technique, attributed threat actor (when known), the affected sector, and the impact and have been detailed in Table 2.1. Using the investigation of these historical attacks, the following key attributes were derived, setting out areas for exploration in the development of synthetic scenarios.

Nation State

As nation-states and organised groups appeared across fifteen of the twenty-three historical attacks from 2009 and onwards, including a level of sophistication into the proposed scenarios accounting for the complexity achievable by a nation-state is of great importance. Their potential ties across historical events demonstrate motivation in the targeting of ICS/OT.

Insider Threat

While attacks conducted by insider threats have become less common in recent years, their ability to allow for the circumvention of security controls and value from a process comprehension perspective (i.e. “the understanding of system characteristics and components responsible for the safe delivery of service (e.g. treatment of water). This includes all relevant physical and computational attributes.”) [91], makes them a significant threat in even the most complex and secure environments. Therefore, accounting for their ability to aid an attack should be considered within the proposed scenarios.

Purely Technical and Socio-Technical

All attacks contain a technical component. This could be the exploration of a system vulnerability, exfiltration of data, etc. However, the prevalence of attacks containing social components (i.e. social engineering) has increased over recent years. In a similar way to insider threats, the exploitation of individuals can be used to circumvent technical controls, particularly a system’s perimeter. Therefore, social vulnerabilities, alongside purely technical vulnerabilities, should be included within the proposed scenarios.

Espionage: Data Exfiltration

Across many of the historical attacks, there exists an element of espionage, that is, to extract useful information. Information acquisition might be the ultimate desired effect of the attacker or act as a precursor to other attacks. There is considerable value to the information held on ICS/OT networks. For example, chemical process information for the manufacture of complex compounds.

Sabotage: Denial of Service (DoS)

While often considered simplistic in execution, the impact of denial of service (DoS) attacks can be significant. As the level of knowledge required to execute a successful DoS is lower than operational process manipulation, it becomes obtainable to a broader range of threat actors (i.e. low and high skilled). Here, sabotage is defined as an observable destructive act that prevents ICS/OT from functioning as intended. Additionally, acts of sabotage have a lower barrier to enact than acts of subversion. Therefore, its inclusion within the proposed scenarios presents an alternate, widely applicable objective.

Subversion: Operational Process Manipulation

Subversion is the act of subtle process disruption (operational process manipulation), which is difficult to detect and may not result in the ultimate destruction of operational equipment. The level of process comprehension required to achieve targeted operational process manipulation is high [91]. However, where historical attacks highlighted the inclusion of nation-state or insider threat actors, the ability to achieve this can be realised. The proposed scenarios should, where possible, also look at options for the manipulation of operational processes by lower-skilled threat actors. This would allow for a broader perspective to be obtained around more strategic, targeted attack objectives.

The baseline requirements derived through the investigation of historical cyber attacks in Chapter 2 form a key starting point in developing synthetic scenarios. However, to ensure they remain valid at a practical level and to better understand their technical construct and execution, they must be developed in a safe/controlled environment. This is discussed in the following section.

3.3.2 Testbed Proof of Concept

Over the last decade, Lancaster University has developed a comprehensive ICS testbed environment [92, 94]. To summarise, it has been constructed through the procurement

Techniques Used	Tools Used
Brute Forcing	Wireshark
Enumeration	Nmap
Exploitation	Metasploit
Lateral Movement	Snap7 (Python Library)
Social Engineering	Custom Scripts
Data Injection	Burp Suite
DoS	
Command and Control	
Device Reconfiguration	
Data Exfiltration	

Table 3.7 Techniques and Tools Used for Scenario Development

and implementation of physical, real-world hardware and software produced by major ICS vendors, including Siemens, Schneider, Allen Bradley, and ABB. This has been leveraged in the physical testing and subsequent construction of the synthetic scenarios outlined in Section 3.3.3, achieved through the practical development and deployment of each attack. This activity was undertaken to solidify the synthetic scenarios' realism further and develop a better understanding of their practical end-to-end execution should it be questioned during the interview process. Table 3.7 presents, at a high-level, the fundamental techniques applied across the development of the synthetic attack scenarios and the tools used to deliver these techniques.

During the practical development of each attack, the devices were explored in use to identify new vulnerabilities which could be exploited to achieve an impact similar to those observed in historical attacks. This resulted in the discovery of two Zero-Day vulnerabilities. These have been appropriately disclosed to the vendors in question. In addition, for ethical reasons, the attack code developed as part of this exercise will not be opensourced.

3.3.3 Synthetic Scenarios

The following diagrams and supporting text provide an overview of the baseline synthetic system architecture, onto which the attack scenarios are applied, including the core operational functionality delivered at a device level. Each attack scenario is broken down into stages, depicted through the use of high-level attack paths.

Baseline Infrastructure

Figure 3.1 presents a simplified ICS baseline infrastructure. This includes a set of devices that are mapped to the Purdue model colour scheme (see Figure 1.1) and are as follows:

- The Sensors and Actuators within the baseline infrastructure are hardwired to the PLC, operating using traditional electronic signalling (i.e. current/voltage based). The use of protocol-based sensors are excluded (e.g. Profibus, WirelessHART, EthernetIP) to simplify this layer of the infrastructure, as it does not form a core component of the attack scenarios.
- The PLC is a Siemens ET200S [249]. This device interacts with the sensors and actuators autonomously through pre-defined control logic and manually via human interaction with the HMI and centralised SCADA system.
- The HMI is a Siemens KTP700F [248]. This device is responsible for the localised monitoring and control of operational processes via the PLC.
- The Remote Terminal Unit (RTU) is a Schneider SCADAPack32 [238]. This device is responsible for collecting and forwarding critical sensor data to the centralised SCADA system from the PLC.
- The Data Historian is a Windows 7 workstation running Kepware [225]. This device is responsible for collecting and forwarding operational data from the PLC to the Data Analytics system.
- The router is a PEPWave [219], and the switch is a Westermo Lynx [282]. These devices are responsible for passing data between each of the Cell/Area Zone devices and the data-centre. While the switch acts passively, the router enforces basic security controls by filtering traffic between its local and remote interfaces. This filtering comes in the form of a rule-set allowing the centralised SCADA and Data Analytics systems to communicate with any device across the Cell/Area Zone. All other communications are blocked.
- The centralised SCADA is a Windows Server 2016 based system running Schneider's ClearSCADA [237]. This application collects and depicts data from the RTU. Its primary purpose is operational alarm generation.
- The data analytics systems operate on Ubuntu Server 18.04, running ThingWorx [121]. This application collects, augments, and depicts data from the Data Historian. Its

primary purpose is the delivery of in-depth analytical and processing capability of operational data.

- The two workstations run Windows 7 and have access to the Data Analytics and Centralised SCADA systems via their web interfaces. One of the two workstations (lighter shade of blue) also operates a client application (ViewX [237]), allowing for a greater level of interaction with the Centralised SCADA system. There are no network-level security controls between these four systems.

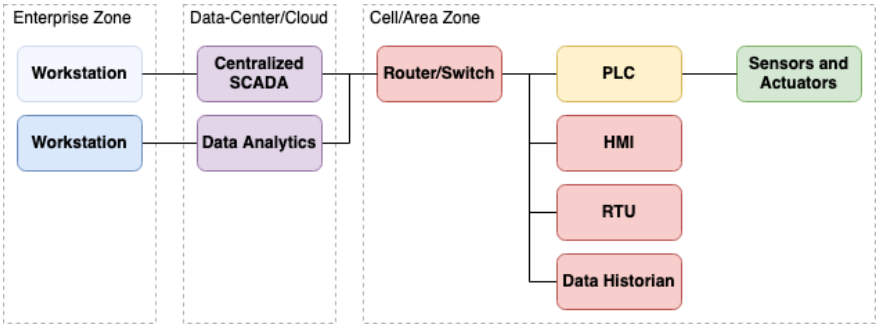

Fig. 3.1 Core Infrastructure

Synthetic Reference Scenario One: Technical - Espionage and Sabotage

The following points describe each stage of attack scenario one. These have been developed and tested within a testbed environment, harnessing tooling described within section 3.3.2:

- Stage 1: Compromise router through the use of password brute-forcing. Once accessed, leverage existing VPN configuration functions and reconnect as a trusted user.
- Stage 2: Enumerate devices on the Cell/Area Zone network. Where possible, extract relevant data (e.g. device configuration and process control logic) for offline analysis.
- Stage 3: Take external monitoring systems offline (i.e. RTU and Data Historian).
- Stage 4: Take the PLC offline.

Synthetic Reference Scenario Two: Socio-Technical - Subversion

The following points describe each stage of attack scenario two. Each stage has been developed and tested within a testbed environment, harnessing tooling described within section 3.3.2:

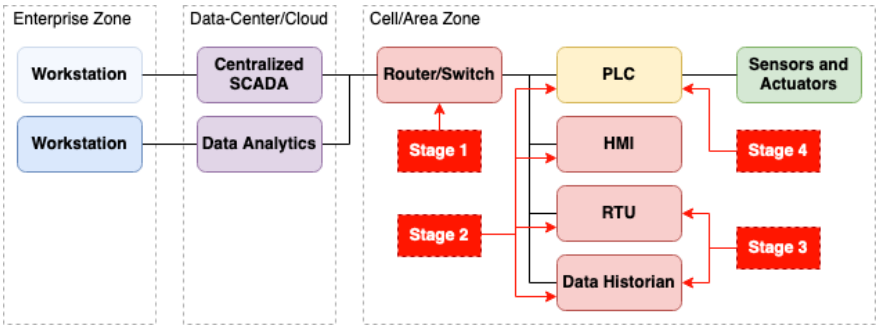

Fig. 3.2 Scenario 1

- Stage 1: Setup an HTTPS listener on a public-facing system. Send an email to a workstation user containing a malicious file. Once opened, an HTTPS session back to the attacker will be established.
- Stage 2: Via the HTTPS session, leverage the compromised user's access to interact with the Centralised SCADA system using its associated client application.
- Stage 3: Via the HTTPS session and access to the Centralised SCADA system, use the inbuilt capability to control the PLC, resulting in operational impact.

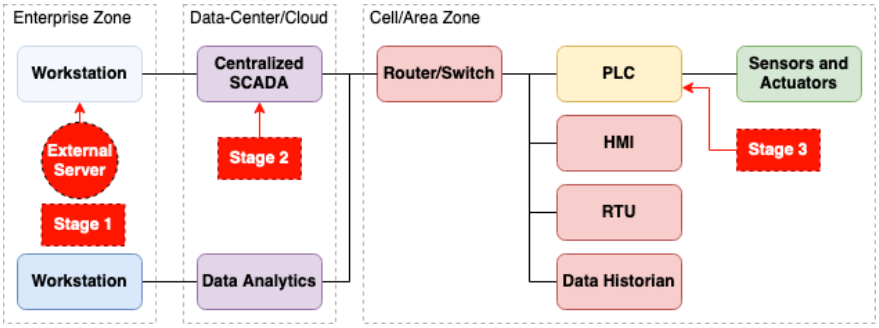

Fig. 3.3 Scenario 2

3.3.4 Summary

The synthetic attack scenarios described across the previous sections were presented at an abstract level. When transforming these for inclusion in the subsequent interviews, the level of detail required may increase supporting participant understanding. In addition, it may be appropriate to provide evidence of the tooling used during testbed proof of concept activities, to further participant's depth of understanding around each stage of an attack's execution.

While the testbed itself has been reviewed to ensure it accurately represents a real-world system through the practical development of each scenario that affords a high degree of confidence in applicability to a real-world context, additional validation was sought through informal engagement with industry experts. Five experts were approached with experience in the field of cyber security consultancy. Each expert was presented with the synthetic scenarios and asked to provide comments on their practical applicability to real-world systems. An explanation was provided on how these had been tested using real-world hardware and software in a testbed environment and how the use of exploits targeting Zero-Day vulnerabilities had been appropriately disclosed. Besides minor changes to the terminology used to describe each scenario, they were accepted as good working examples of attacks that could be executed against an ICS/OT environment.

The following section describes how the developed synthetic scenarios have been utilised in a series of interview with industry stakeholders to stimulate discussion on cyber incident response and recovery processes and the use of standards and guidelines in practice; discussed in Section 3.2.

3.4 Stakeholder Interviews

The following subsections provide an overview of an empirical study with industry stakeholders operating/regulating elements of European CNI. Using semi-structured interviews and synthetic cyber attack scenarios, this study explores current cyber incident response and recovery practices, the adoption of existing standards and guidelines, and challenges in their use. The goal of these interviews is to further investigate the findings from Section 3.2 and to assess whether the challenges from using standards and guidelines are observable in practice.

3.4.1 Methodology

In search of appropriate research techniques, interviewing key stakeholders working across the topic area was applied. A simple ethnographic observation would prove extremely challenging and time-consuming, particularly when considering the sensitive nature of systems and processes being evaluated and the requirement to seek approval from both participants and the organisation in which they work. Interviewing was selected as an appropriate alternative [216], enabling each participant to discuss response and recovery activities without direct reference to a specific organisation or system. The ability to explore meanings, routines, behaviours, etc. [231] all adds towards appropriate focusing, particularly when

discussing complex topics, and confirmation of meaning from both parties (the interviewer and interviewee) may also be required [30].

Interviewing typically falls within three core categories, viz. structured, semi-structured, and unstructured. Here a semi-structured approach has been adopted, often seen as the most common form of qualitative research methods. This approach provides adequate flexibility with a pre-defined core question set, options to include improvised follow-up questions, and explore meanings should they be required [13]. Where existing cyber incident response and recovery guidance discussed across Section 3.2 highlighted deviations in provided detail, the possibility of a repeat scenario was considered in the selection of an interviewing technique. The flexibility offered through a semi-structured approach presents significant benefits, allowing for additional probing where little detail is provided and further exploration of more comprehensive approaches where required. The following sub-sections break down points considered through the application of this approach.

Sample

In selecting an appropriate participant sample, the aim is to understand the topic area from all relevant perspectives. To achieve this, a broad approach to the targeting of participants was applied. This resulted in a diverse collection of role-profiles. More specifically, those engaging in cyber incident response and recovery processes across multiple systems, with varying levels of responsibility. This sampling approach provides multiple perspectives, building a broader picture of how cyber incident response and recovery activities are conducted.

To summarise, eight participants were selected holding the following roles:

- Chief Information Security Office (x2)
- Operational Technology Manager
- Information Technology Manager
- Information Assurance Manager
- Engineering Delivery Manager
- Emergency Arrangements Coordinator
- Operational Technology Software Engineer
- Operational Technology Cyber Security Inspector

The levels of experience varied amongst participants within each of the defined roles. The majority of which, however, had been working with industrial systems for over ten years.

Threats to Validity

The work of Campbell & Stanley [29] discusses common threats to the validity of the collected data. To address these issues, attention was initially focused on the participant sample, as previously discussed. From this, interviewing techniques applied to build rapport, trust, and openness were adopted, with questions covering all relevant topics and those topics alone (i.e., no irrelevant questioning). As this set of interviews is designed to complete orientation on cyber incident response and recovery, included questions were drawn from initial understandings achieved through the analysis of standards and guidelines in Section 3.2.

Where Powney & Watts [223] discuss the emergence of interesting content upon completion of interviews (i.e., when the recorder is switched off), notes were taken and added to the interview protocol/guide. Subsequently, additional prompts were included for potential re-interviews of the same interviewee, for inclusion within other interviews, or simply as salient content worth exploring as part of further focusing efforts. This approach allowed the interview process to evolve in a structured and managed way while eliciting pertinent information.

Reliability

Of primary concern to the reliability of collected data is interviewer bias. This is the ability to trust that findings are not derived from research instruments or as a result of an interviewer's quirks and improvisations. Concerning this is the perspective of "insiders". Insiders can be defined as interviewers who share similar cultural, ethnic, linguistic, national, and religious heritage to interviewees [82]. More simply, where the interviewer and interviewee are part of the same organisation (i.e., work colleagues) [13]. This can prove highly valuable when seeking additional participants, understanding organisational structures, etc. [13]. However, it can also increase the risk of data reliability issues, with a higher probability of assumptions and general interviewer bias, based on an interviewer's perspective of "the way things are". While the research team may or may not be considered insiders by these standardised definitions, having collectively spent over thirty years working for CNI operators, each project member may have their own perspective on "the way things are", from organisational culture to policies, power relationships, etc. The positive attributes of these experiences were utilised in the interview protocol/guide design. However, to account for the possibility of negative attribute inclusion, this guide was read and understood by all project members prior to the start of interviews.

Rubin & Rubin [231] refers to transparency and consistency; this accounts for consistency and evidence of inevitable inconsistencies in data. These were appropriately handled and included within the analysis phase.

Neutrality beyond the aforementioned "insider" bias was also considered throughout the interview protocol/guide design process and during each interview. As an interviewer, acknowledgement of personal background, age, class, etc., can all influence an interview's direction and output.

Primary Practical Technique (In-Person Interviewing)

In-person interviewing provided the mechanism for engagement, as this allowed for clear and in-depth data collection. This interviewing technique provides additional information compared to remote interviewing, mainly due to facial expressions and visual cues. In-person interviews can also be considerably longer than remote interviews since participants have provided greater commitment to participate and are less likely to be distracted during the interview process [13].

Interview Protocol/Guide

Each interview was broken down into the following six stages, providing a logical structure to the interview protocol/guide:

- Preface
- Establishing Demographics
- Scenario Familiarization
- Response and Recovery Analysis
- Guidance Analysis
- Conclusion

The core focus of these interviews was to build upon Section 3.2, providing a broader understanding of cyber incident response and recovery practices within an ICS/OT context. More specifically, how key stakeholders broach cyber incidents. Taking direction from Section 3.2, the questions aligned to these interview stages are aided through the inclusion of probes and definitions. Due to time limitations, additional probes were only used to provide a greater depth of understanding to directly-related, salient points of discussion. The following

provides a summary of primary interview questions. The complete protocol/guide can be found in Appendix A.

Establishing Demographics

The following question-set was applied to the demographics phase.

- Please can you tell us your job title and provide a very brief overview of your core roles and responsibilities?
- How many years of experience do you have working in this role?
- At a high level, please can you explain to us what you understand the term Response and Recovery to mean within the context of an Operational Technology (Industrial Control Systems) cyber security incident?

Response and Recovery Analysis

The following question-set was applied during the response and recovery analysis phase, once the participant had been shown the first synthetic cyber attack scenario. Upon completing these questions, the participant was then shown the second scenario and asked if anything would be done differently.

- Given your role in the organisation, at a high level, what are the core steps you would go through as part of response and recovery operations in the example scenario?
- How many individuals within the organisation would work directly with you on these steps, so performing the same role as you?
- Who else would you have direct engagement with during response and recovery processes?
- How many individuals across the organisation would be involved in response and recovery operations more generally speaking?
- When undertaking a response and recovery operation to this scenario, what do you consider the primary goal to be?
- When you are undertaking individual response and recovery actions, how do you factor in risk evaluation as part of the decision-making process?
- Typically, what are the expected outputs post-incident, once you have appropriately recovered from an incident and everything is back to normal?

Guidance Analysis

The following question-set was applied during the external guidance analysis phase.

- In your opinion, which standards or guidelines best cover response and recovery in relation to Operational Technology cyber-attacks?
- As a final question, what is your opinion on currently available standards and guidelines within the context of cyber incident response and recovery?

Conclude
The following question was applied during the conclusion phase.

- Would you like to add anything which may be relevant?

Analysis

In search of an appropriate methodology by which captured interview data could be analysed, template analysis was selected. Also referred to as "codebook analysis" and "thematic coding", template analysis offers a highly flexible method to the analysis of qualitative data [145]. Sitting between the relatively rigid approach of content analysis in which analytical codes are all pre-defined [229], and the opposite approach of grounded theory in which all analytical codes must be derived from the data [88]. This approach was initially conceived by Crabtree & Miller [45], and was later adopted by King et al., [145], from which it saw an increase in adoption across a variety of fields. Considering participant numbers and their diverse roles, the flexibility offered through template analysis provided significant value over alternative approaches, allowing to create an initial code-set aligned to core areas of interest, with relevant additional codes added as they emerged.

While template analysis has fewer specified procedures, offering greater flexibility to statistical and qualitative analysis of the same data, recommendations are proposed by King et al., [145], these were followed within the analysis of interview data here. For example, through the use of the previously described interview protocol/guide, an initial code-set was constructed but was limited to allow for further granularity or abstraction if required. Where too many pre-defined codes may constrain/confuse analysis, too few may cause a lack of direction. Undertaking a brief review of initial transcripts allowed additional codes to be added. This helped build confidence in the code-set before starting the complete data codification.

3.4.2 Interview Results

Key findings from the interviews with stakeholders are summarised here. These have been grouped based on identified themes and key points of interest. It is worth emphasising that all of the points discussed here have been identified from the contents of the interviews.

These, therefore, reflect the generalised opinions of participants but may differ from person to person.

Primary Goal

Before identifying core phases individuals enact during the cyber incident response and recovery process, it is important to understand what participants believe the primary goal should be. As one would expect, a fundamental focus on safety through incident containment and the preservation of critical system integrity rose to the forefront, with a couple of additional notable points raised once the system was safe.

'The primary goal for us on site would be to make it safe.'

Preserving evidence allowing for subsequent investigations to better understand how the attack happened was raised by many participants in varying forms. This was also tied to appropriate chains of custody.

'Preserving evidence so that we can ensure then that we don't lose how it happened.'

Establishing communications between all relevant parties as quickly as possible was deemed highly important, from the person who detected an incident to shift managers and beyond. From this point onwards, participates felt entire pre-prepared response and recovery arrangements would fall into place.

Those in technical roles followed up on these goals with practical actions based on the two hypothetical scenarios. For example, removing all external communications points to prevent the continued manipulation of systems.

An emphasis was made by one participant on the balance that was required between providing safety, security, and business continuity and that a risk-based decision needed to be made based on the type of incident that was being responded to.

'The following questions need to be asked: do I continue operation? Is there a safety issue? Or is there a security issue where information or material and the protection of it can be undermined?'

Key Phases/Tasks

As a reminder, the identified phases/tasks are a direct output of discussion formulated through the developed hypothetical scenarios.

'... if it comes from operators, the first response is probably "something's broken" as opposed to "something's been hacked"... '

Incident identification could come from several parties. For example, Security Operations Centres (SOCs) may have detected suspicious behaviour based on various data capture points. Alternatively, operators using impacted processes may be the first to highlight suspicious activity due to a loss of control. In the instance where operators are the first to raise concerns, it would likely be viewed as a system error instead of a cyber attack.

'Our arrangement would require the team that discovered the breach to contact our 24/7 site shift manager... '

Upon a system issue/incident being raised, escalation processes would be enacted. These typically involve notifying senior site managers in the first instance. A central incident control team would be formed, including key personnel to support decision-making processes. As scenario one is diverse, spanning both IT and OT systems, forming a centralised response team would include individuals from both sides of the organisation. In addition, specialised external support would be brought in. This support could come in the form of cyber security specialists, forensic analysis from third-party companies, for example. Alternatively, it might be operational engineering or safety-focused resources from partner or parents organisations.

'... he would need to know a high-level summary: what does it mean to the site, what could the impacts be, how long until you get it fixed, is it going to spread, etc.'

As a collective, the central incident control team would provide a diverse skill-set and knowledge-based to those in charge. Therefore, translating technical information into a language all parties can understand regarding the capabilities of the threat and what the potential impact could be, was deemed of great importance. This would provide a platform for response and recovery decision making moving forwards, including the initial formation and continuing involvement of appropriate personnel.

'We encourage the reasoning for decision making to be including on these log pads.'

Participants referred to priority systems, ticketing systems, logging systems, etc., as part of the fundamental setup of centralised incident control centres. These systems are brought in to allow for a coherent understanding, management, and recording of all subsequent response and recovery actions. This includes the criticality of the incident, reasoning behind decision making, and timings of actions taken and their resulting output.

'First of all, get it safe and secure. Preserve the evidence. And then look for how it happened with a view to then prevent it from happening again.'

Regarding the general operation of systems under attack, several factors came into focus aligned to the continued running vs shutdown of processes. These were naturally driven by safety; however, forensic data to support follow-up investigations was considered beyond this. The value of evidence was viewed as critical to better understand how the attack took place, where the gaps in system defences lie, and thus where additional focus is required moving forwards. This was also considered critical to subsequent recovery processes, giving confidence to the execution of follow-up actions, e.g. reloading backup configuration and knowing that will have the desired effect. Furthermore, the criticality of specific systems to the success of the business may be high. This could again deter from a complete shutdown where there is no risk from a safety perspective.

'We'd be looking for things on the intrusion prevention system and we'd be engaged directly with the SOC team to identify whether they've identified malicious activity.'

Response documents would be used to guide decision making throughout an incident. However, based on the hypothetical scenarios posed to participants, several key actions were noted from a technical perspective. These include, but are not limited to:

- Examine operational system logs.
- Examine intrusion prevention system and firewall logs.
- Examine web-proxy logs.
- Engage with SOC teams.
- Engage with forensic teams/conduct an analysis.
- Identify impacted workstations.
- Isolate impacted workstations.
- Block access to the attackers IP/Domain.
- Remove external connectivity.
- Check systems are up to date (e.g., OS and AV).
- Erase/replace devices and restore from backups.

These actions can all be used to support decision-making within the centralised incident control centre and would be communicated promptly to ensure all parties are kept up-to-date with each team's overall process. Thus, providing a complete picture of the incident and add assurance to the actions taken and their ability to restore normal, safe operations.

'The investigations are generally based on how wide an impact another incident like that could have. It starts at purely local to that plant. Then it's across the site as a whole. The highest level is anything that could happen to other sites.'

Considering post-incident outcomes, several relevant points have been raised throughout this section. However, a more comprehensive analysis of an incident would be undertaken. Differing levels of analysis and investigation would occur based on the impact of an incident or potential impact of a similar incident in the future.

'After every incident we do a lessons learnt and implement actions on how to improve the resilience of our systems. We raise what we call a learning report and then a manager is assigned to that and then actions are put in.'

As a starting point, all logs taken during an incident would be reviewed. These, along with forensic and more technical details, can be used to form the basis of a lessons learnt/incident report. In addition to lessons learnt for internal review, these sources could also be included within any legal challenge or regulator inquiry, providing a clear timeline of events and rationale for every decision that was made. The report would typically provide an executive summary and an action-based output to address the root cause, offering suggested additional measures that could be applied to prevent any reoccurrence.

'There's no point in recovering and then being in the exact same position and them doing the same thing.'

One participant explained that he would expect dutyholders to follow a playbook to take appropriate response actions in a timely manner.

'A lot of our dutyholders have playbooks to contrast which behaviours are anomalous and help them identify what is going wrong.'

A point was also made on how response actions could have adverse effects on evidence collection, highlighting that if a timely response was not critical, delaying the recovery process in favour of collecting evidence was recommended.

'There would have to be a decision made based on how much time can be dedicated to evidence gathering which would be A) to conduct a criminal or internal investigation and B) to inform Lessons Learnt since we would need to turn on the industrial zone and get it moving again as quickly as possible.'

Risk Evaluation

During response and recovery processes, it was widely acknowledged that important decisions must be made quickly to ensure systems remain stable, ensuring no safety-related risks are realised. Several interesting points were raised concerning the evaluation of risk during decision making, from risk owners to formal and semi-formal processes. However, it is important to note that risk evaluation was deemed cause agnostic. Therefore, the following points are mainly applicable to incidents caused through any means, malicious or otherwise.

'It's the job of the shift manager in the operational centre and the job of the controller of the day in the tactical centre to surround themselves with a team of specialists who provide advice to them.'

Starting with those involved in crucial risk decision making roles, while central teams are formed to manage an ongoing incident, they rely on subject matter expert knowledge from across the organisation. For example, each operating system may have a supervisor; this role is ultimately responsible for ensuring its safe operation. As a result, supervisors will have an in-depth knowledge of operational processes, response documentation, the potential impact given actions could instigate, etc. Therefore, they would be called on to provide advice to central teams throughout an incident.

'We would have a response document and follow the right protocol.'

The processes for evaluating risk were mixed, formulated around understanding the consequences of all possible actions. Formal risk evaluation procedures did exist for specific scenarios, in addition to response documents with defined protocols. However, evaluation techniques, on the whole, were described mainly as semi-formal decisions made on the collective agreement of key stakeholders (as previously described), with the operational centre acting as a primary responsible party. However, each operational process's pre-existing assessments indicate their criticality and the impact that could be realised during an incident. Should an incident escalate to the point where individuals would be required to enter potentially hazardous areas of a facility, rigorous formal processes would be followed.

'There's a formal process including a briefing, agreement, and sign off, at different levels of authority before we'd send someone into a potentially hazardous situation.'

Safety formed the key consideration of risk evaluations. This was supported by environmental factors, the integrity of forensic resources/maintaining the chain of custody, and finally, reputational damage. Adding a reputational viewpoint was described as an attribute that could impact public perception, stimulating negative rhetoric on facilities' continued operation, and therefore should be considered where possible.

'The consequence of the risk would be based on safety first including environmental factors, and then reputation.'

As previously noted, pre-existing risk assessment documentation is used during decision-making processes. In addition to these, emergency arrangement teams regularly exercise incident response processes, refining them to pre-emptively capture all risks that could arise during an incident and improve the organisations ability to effectively minimise impact.

'Our emergency arrangements team would have exercises on this sort of thing to improve it and test it.'

Risk-based decisions would also change based on circumstances which could affect the impact of specific actions. For example, within the energy sector, an abundance of power is available during the summer as opposed to winter. The amount of risk associated with shutting down a power plant would, therefore, differ depending on the time of year.

Human Capital

Considering the human capital required during an incident, the previous sections have offered an initial insight. However, it is an important point for which additional detail is required.

'If I wasn't on duty, my line manager would take on my role as he is the senior C and I engineer.'

Considering the role profiles of participants, there existed very few examples of identical roles. However, each organisation's structure was constructed to account for loss of coverage, i.e., although in differing roles, individuals could step into their colleagues' shoes and perform the role required both up and down-stream.

'If there's an emergency on site, we always say that the whole site is at our disposal. Everyone understands that there's a responsibility to do what they need to do.'

Taking a more holistic view, it was considered that every individual working on a site would be at the disposal of the central emergency response team during an incident, with 24-hour on-call personnel covering key role profiles.

'We also have a contract with an external company. To do things like investigations and forensics. . . '

Extending out from localised on-site resources, all participants raised the possibility of pulling in additional resources from specialised third-party organisations (e.g., cyber security practitioners and forensic investigator), in addition to operational and engineering personnel from across the wider business, partner, or parent organisations, dependant upon the skill-set required.

'I would probably become more of a support to the OT side of things....My role would shift a little bit to lace the forefront.'

The aforementioned measure makes it hard to define precise numbers or skill-sets required during an incident, particularly given the scenarios' diverse nature. However, it was noted that specific roles could act in more of an advisory capacity where an incident has occurred on systems outside of their control or direct expertise.

If an incident were critical enough to affect multiple sectors or bring about danger to civilian life, then a range of stakeholders in the private sector, the regulatory body and the government would convene and decide upon the most appropriate action to take.

'Legislation and cooperation between stakeholders would allow timely decisions to be made to ensure that a safe outcome would be achieved.'

Use of Standards and Guidelines

Opinions of existing standards and guidelines were mixed amongst the participants. Some considered them to have matured over recent years, provide a valuable base of expertise. Others believed they were too focused on one specific domain and failed to capture similarities between IT and OT, resulting in inefficient activities.

The prevalence of existing standards and guidelines for use during the planning and execution of response and recovery activities was limited. Throughout the interviews, examples were discussed mainly from a higher-level viewpoint. This was achieved by identifying specific organisations responsible for standard and guideline development, as opposed to individual resources (e.g., The National Institute of Standards and Technology (NIST), not NIST SP800-82, SP800-53, etc.).

One participant also discussed an internal framework currently under development by centralised teams within their organisation, incorporating cyber incident response and recovery. However, the participant could not advise on its creation at a technical level and the level of concept/methodological coverage from existing standards and guidelines beyond an alignment to NIST.

No approaches derived from academic works were discussed whereas existing standards and guidelines were discussed positively, highlighted attributes focused on an increased drive towards their adoption and increasing levels of maturity.

'. . . awareness was very limited prior to 2011-2012.'

The awareness level of appropriate standards and guidelines was noted as limited before 2011-2012. However, in parallel to the increased awareness over recent years, the level of standard and guideline maturity is also believed to have developed, increasing its value to this new audience.

'What's the point in us reinventing stuff when someone's already done it?'

Through the use of internal reviews, participants felt they had demonstrated good coverage across multiple aspects of cyber incident response and recovery. However, acknowledgement was made towards the use of existing expertise through the adoption of documented approaches rather than reinventing from the ground up.

'Historically, cyber was just considered to be "for the IT guys"; it's not anything to do with us.'

A broader acknowledgement of cyber security and related standards amongst non-IT-based personnel demonstrates an increased maturity level from an operational perspective. This involved conducting reviews as a collective group from multiple business areas to better understand requirements and the importance of cyber security from the more common health and safety viewpoint.

'It's not something that we could just say that "it's the geek stuff, you sort it out". . . I align it very much with an important health and safety focus.'

Where existing standards and guidelines were viewed in a negative light, a variety of points were raised. These were aligned mainly to variety, length, applicability, and complexity.

'There isn't particularly one that's the "magic" one'. . . '

Initial insights identified that no single resource was deemed appropriate for all aspects of cyber incident response and recovery and that, as a result, internally developed sector-specific approaches were under development using existing standards and guidelines as a base.

'…they're very IT focused and therefore very information focused as opposed to function.'

Participants with an OT background believe that existing standards and guidelines often lacked tooling and frameworks to adequately cover OT systems, applying an IT focus based on information instead of function. These were seen as barriers to their adoption.

'We'll have some guidance that is created for IT and then we'll end up having to do exactly the same work but created for OT. There's a lot of parallels and similarities.'

In contrast, participants with an IT background raised questions around the requirement for independent guidance, stated similarities exist between concepts from both IT and OT domains. This leads to a feeling of continued isolation, resulting in a counter-productive use of time and resources.

'I think that better guidance on how to implement said guidance and best practices would definitely help take the pain out of it all.'

The volume and depth of existing standards and guidelines from a usability perspective raised concerns about their application. With resources stretched, the ability to explore, understand and implement existing standards and guidelines dramatically increase the barrier to entry. This could result in inconsistencies, with participants demonstrating a desire to know everyone is efficiently pulling in the same direction, with approaches suitable for their organisation's size and scale/risk profile.

'…some of the guidance goes into too much detail rather than what we actually need.'

In more general discussions with key cyber security personnel, it was acknowledged that non-cyber security-focused colleagues would have a lower awareness level, but thanks to broader work programmes, it was now at a higher level of maturity than in previous years. In addition, processes and time allocation for reviewing existing standards and guidelines had been established.

One participant argued that exercising and training, which can be derived from standards and guidelines, provides a significant benefit to dutyholders as it pushes them to be more familiar with processes thanks to hands-on experiences.

'As an engineer, what appears to be the best way is to get groups of people together and exercising them to develop muscle memory that can be used during real incidents.'

3.4.3 Summary

Across the previous sections, the methodology applied to a set of interviews with individuals working in and around ICS from both IT and OT backgrounds has been outlined. This included a pre-defined question-set, allowing for a degree of flexibility through semi-structured in-person interviews. The output of which was analysed using template analysis, a suitable technique given the nature of the research objectives.

During each interview, several themes emerged, covering key topics from existing response and recovery practices to the level of internal and external personnel engagement and opinions/use of existing standards and guidelines within a response and recovery context. These, along with findings from the initial analysis of existing standards and guidelines in Section 3.2, will be discussed in more detail across the following section.

3.5 Discussion

The following discussion is broken down into existing standards and guidelines and engagement with industry stakeholders. These two studies, the former focusing on the theory behind incident response and recovery and the latter focusing on its implementation in practice, offer input into improving ICS/OT cyber incident response and recovery capabilities. The goal of these studies was to identify the challenges faced when using standards and guidelines documents to improve and/or assess ICS/OT cyber incident response and recovery capabilities. These findings are summarised here.

3.5.1 Guidance

The analysis of existing guidance across Section 3.2 captured thirty-one resources in total. This was made up of both UK and International standards and guidelines from governmental organisations (NCSC, NIST, HSE, DWI, NRC, CNSC and ANSSI); non-statutory organisations (ONR and NERC); international organisations (ISO/IEC, ENISA and IAEA); educational institutions (Carnegie Mellon University and SANS); and industry institutions (NEI and CREST).

This vast array of material demonstrates an abundance of guidance ICS/OT operators can consult towards developing their own internal processes and overall capability. Furthermore, it was found that these resources are often interwoven with one another, acting as key multi-directional reference points.

The analysis of the thirty-one identified resources found a lack of consistency in the breadth of content when aligned to a holistic criteria set (See Tables 3.4, 3.5, and 3.6). While

consulting a single resource could lead operators to review multiple additional cited resources, this may not always be possible. Paywalls, for example, can impact accessibility to cited resources. Furthermore, where baseline information is included around a specific cyber incident response and recovery phase, it may be misunderstood as complete, with additional cited materials considered optional.

The adoption of processes supporting cyber incident response and recovery can be both technical and/or procedural in nature. While guidance must adapt to its intended audience (e.g. non-technical managerial positions versus engineer-level security specialists), it is also vital that topics are covered at an appropriate level of detail to enact meaningful paths of progression. In reviewing existing resources for technical versus non-technical content, several instances were found where the required level of technical detail was limited or not present. The value of non-technical discussion was acknowledged in conveying critical concepts; however, implementation can be challenging without supporting technical specifications and direction.

A wealth of information can be found across the resources reviewed here. However, the isolated selection of a single resource to drive change within an organisation will likely result in a less than complete picture. The quantity of available resources also presents a challenge for operators. How does an operator know they have selected the most comprehensive resource or set of resources? Beyond regulatory interaction, how does an operator know they have implemented cyber incident response and recovery processes at an appropriate level of technical depth? Without a clear overview and understanding of a broad resource pool, as provided here, answering these questions can present a significant challenge.

3.5.2 Stakeholder Engagement

During the initial demographic question base, it was established that most participants had only ever worked in one industrial sector. Career opportunities to develop pathways into specific technical and managerial roles were commonplace. The in-house/in-sector development of personnel is logical; however, it can lead to isolated viewpoints without external engagement. While external engagement can be a challenge due to the justifiably closed nature of operational facilities, engagement with relevant third parties can prove to be highly valuable when developing holistic cyber security capability.

Some participants described the in-house development of tailored cyber incident response and recovery approaches. It is unclear on the level of external engagement being undertaken to obtain a third-party viewpoint. However, it was noted by several participants that reinventing the wheel is undesirable, and taking input from existing standards and guidelines is a preferred

approach, with internally developed approaches using well-known materials (e.g. from NIST). This provides confidence and credibility to the development of tailored internal guidance.

The processes outline towards a central incident response team's formation, and operations were well understood by all participants. The ability to leverage all internal, and bespoke external resources where necessary, appeared almost limitless, with contracts in place to support every eventuality. Given the cause-agnostic nature of central incident response team processes, a clear understanding of procedures/requirements allowed for a smooth, well-orchestrated establishment process. The level of internal resources to support cyber incidents from an OT perspective was unclear with all participants; this could cause delays in identifying an attack's progression and maturity but would not cause significant challenges in reacquiring control and, therefore, the integrity of systems from a safety perspective.

During response and recovery activities, the documentation of system state, decision-making processes, actions, and their subsequent effect, were well described by all participants. The value of documenting actions during an incident was clearly articulated, from future use during legal or regularity challenges, to root cause analysis/the technical understanding of how an event occurred. Having such a comprehensive approach supports not only an understating of how something happened but what can be done to mitigate a similar event occurring in the future and what decisions helped/hindered response and recovery efforts. Findings of this nature can be fed into future hypothetical exercises and overarching processes to enhance skill-sets and an organisation's overall ability to effectively respond and recover to previously unseen incidents.

When considering the evaluation of risk during response and recovery decision making, a semi-formal approach based on input from a broad range of experts was adopted. While this was focused mainly on the implication actions could have on safety, they also considered environmental impact, forensic data integrity, and reputational damage. Formal evaluation techniques were applied to specific scenarios, where a situation dictates a requirement for personnel to enter potentially hazardous areas, for example. However, it was deemed impractical in a time-critical situation to cover every eventuality, thus opting for a semi-formal, cause agnostic, expert input-based approach.

There existed some conflicting views on standards and guidelines, with IT-focused participants stating that they could see direct similarities between IT and OT tailored resources, whereas OT focused participants believed them to be too information focused (as opposed to function-focused), their value was considered significant towards maturing existing cyber security processes. This was echoed throughout with a desire to take existing, proven approaches rather than reinvent them from the ground up.

Lessons learnt from an OT cyber security perspective appeared less mature than other areas. This is unsurprising due to its relatively recent formation when compared with conventional engineering and safety-focused cases. The involvement of individuals with a broad range of skill-sets within central incident response teams, and subsequent follow-up lessons learnt, is currently the closest way to comprehend OT focused aspects, with input from security and engineering personnel. The use of lessons learnt reports within cyber exercising could also be seen as a pathway to the overall development and understanding of cyber security challenges across an organisation.

The engagement from participants in internal and national-level cyber incident exercising can be seen as a positive step in developing capability and overall preparedness. Although some operators are mandated to perform exercising, some of the participants engaged voluntarily. This commitment forms the most practical route to test new cyber-focused response and recovery practices, whether derived from standards and guidelines or lessons learnt.

3.6 Conclusion

This chapter provided a window into cyber security incident response and recovery guidance, alongside a high-level overview of processes adopted by operators. In extending the scope of Section 3.4 to capture opinions on existing guidance, an understanding of how they are currently viewed and used in practice has been provided.

While significant effort has been invested by reputable organisations in the creation and continued evolution of guidance to support operators develop cyber security incident response and recovery capabilities, its uptake could be improved. The volume of guidance, its intertwined nature, and varying levels of scope present challenges in its adoption. Selecting a guidance set that provides only high-level non-technical information, coupled with the exclusion of supplementary cited materials and limited coverage across the defined criteria-set, could leave operators lacking core skills, implementing under-developed supporting technologies, and operating limited overarching processes.

When considering the internal growth of talent within industrial organisations, it becomes critical to provide comprehensive guidance allowing for new roles and career paths to form. The use of existing standards and guidelines to develop internal processes provides a primary conduit towards identifying required skills, and general human capital, further highlighting their importance.

As all personnel can be used during an incident, including the involvement of third parties, through an enhanced cyber security understanding, gaps in human capital may be

identified. For example, during exercising, an organisation will be able to identify that while key response and recovery phases would significantly benefit decision-making processes within central incident teams, increasing the efficiency of activities, preserve forensic data, etc., they require additional personnel to be recruited, or existing personnel to undergo additional training.

Based on these findings, providing a framework that could be used to identify/assess an organisation's existing overarching cyber security incident response and recovery process coverage would directly benefit operators. Where gaps are identified in existing practices, it becomes vital to better understand how they can be developed. In the interest of avoiding the recreation of existing material, a framework's coverage of response and recovery phases should be based on those detailed in existing standards and guidelines. Furthermore, specific section numbers from within each of the references standards and guideline should be highlighted to avoid the requirements for a comprehensive and resource-heavy review by each framework user. A framework of this nature would complement existing processes, offering a high degree of credibility, instilling confidence in its use.

Initial framework concepts were discussed with interview participants and received a positive response. Therefore, the following chapter introduces the proposed framework, acting as a starting point towards supporting operators in developing their cyber incident response and recovery capabilities.

Chapter 4

The Industrial Control System Cyber Incident Response and Recovery (ICSCIR&R) Framework

The following cyber incident response and recovery framework has been created based on the findings of the two subsequent studies discussed across Chapter 3. From these studies, two key points were identified: firstly, existing guidance lacks consistency in the breadth and depth of information provided, and, secondly, a single resource by which existing cyber incident response and recovery processes could be reviewed for completeness and further developed, would offer significant value to ICS/OT operators. In the interest of avoiding the recreation of existing material, an undesirable option raised during the stakeholder engagement, the framework presented here focuses on aggregating information across the previously investigated thirty-one international standards and guidelines. The core output of which provides a centralized, credible resource used to review, support, and enhance an organization's cyber incident response and recovery capabilities.

4.1 The ICSCIR&R Framework

From the analysis of thirty-one international standards and guidelines in Chapter 3, Section 3.2, four high-level cyber incident response and recovery phases, aligned to seventeen sub-phases, were identified. These were summarized in Table 3.4, and are used as a base for expansion in the framework. the second study, the stakeholder engagements, discussed in Chapter 3, Section 3.4, provided an in-depth practical understanding for developing the process-flow and contents of the framework. Due to the size of this framework (See

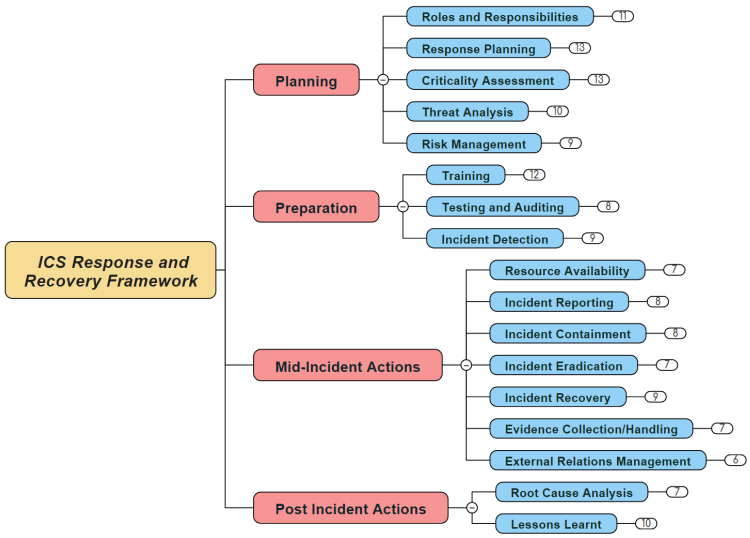

Fig. 4.1 Cyber Incident Response and Recovery Framework [152]

Figures 4.1 and 4.2), it has been turned it into an interactive HTML resource available on Github [152]. The following subsections provide a breakdown of the information aligned to each sub-phase within the framework and its overarching modes of operation.

4.1.1 Overview

A high-level description of each sub-phase, allowing framework users to view their core functions. This helps in the selection of sub-phases for further development.

4.1.2 Dependencies

While each sub-phase has its own unique set of outputs, those outputs may feed directly into subsequent sub-phases as pre-requisites. The high-level view of such dependencies ensures framework users account for sub-phase interplays.

4.1.3 Example Checklist

Sub-phases can be highly detailed, taking time to understand and develop. However, the inclusion of example checklists offers an initial starting point for framework users to explore their existing capability.

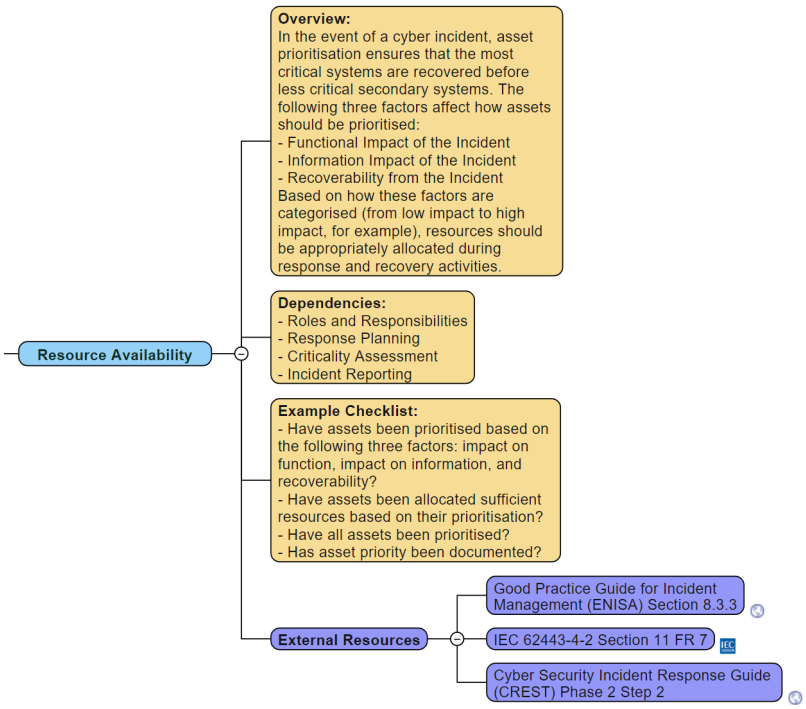

Fig. 4.2 Resource Availability sub-phase of R&R Framework [152]

4.1.4 Additional Resources

Providing the most critical element of the framework are additional resources. Here, each sub-phase is mapped to specific sections of the thirty-one standards and guidelines, saving framework users time in their inception/continued development.

4.2 Framework Operation

The framework provides a light-weight, highly accessible resource that can be used in two primary ways: (1) to review each of the identified cyber incident response and recovery phases/sub-phases against existing capabilities, supporting the identification and development of existing gaps, and (2) as a quick reference guide to understand, assess, and develop, specific phases/sub-phases.

While the framework does not provide a quantifiable methodology towards assessing existing capability, its use as defined within this section offers a high-level, flexible approach to identify gaps and deficiencies in existing practices. More importantly, it provides highly-focused direction to credible resources allowing for the continued development of cyber incident response and recovery capability. These resources provide guidance on the creation of policies and process, including those directly associated with practical security controls,

affording framework users not only with comprehensive scoping coverage but depth in practical detail. The framework can be of significant added value to CNI operators in carrying out their everyday tasks and can guide them and their managers in selecting the suitable implementation for their environment. As such, to not be prescriptive, the framework is defined as a means to guide operators towards the appropriate tools rather than to define specific rules and processes.

Figure 4.3 has been created to support user understanding of the framework's operation. This figure depicts a process flow aligned to the two primary methods of use. The first of which would see a cyclic flow from the initial cyber incident response and recovery phase (Roles and Responsibilities) to the last (Lessons Learnt), whereas the second would involve a single pass on the relevant sub-phase of particular interest to the user (i.e. to further develop know issues in related current practices). The stages of this process flow are as follows:

- To begin, the relevant cyber incident response and recovery sub-phase should be selected from the framework using its associated title. This action can be supported by using the high-level overview, included as part of each sub-phases supporting text. Where an initial sub-phase has been identified but does not match the user's requirements (a possibility with the second method of framework use), a step back to re-review alternative sub-phases will be required.

- Using the provided checklist aligned to the sub-phase under review, the user should conduct an assessment of current capabilities. This activity provides a high-level view of current capabilities vs sub-phase requirements and acts as a starting point to better understand the associated sub-phase and whether it has been considered within existing cyber incident response and recovery processes.

- From the initial baseline checklist, associated dependencies should be reviewed. This begins to build a more comprehensive picture of the sub-phase under review, its key characteristic, and how it fits within the broader cyber incident response and recovery life-cycle. If they are met and understood, no further action is required. Alternatively, a loop back to review each dependency within the framework could be conducted.

- Where the information provided within the framework is sufficient, the sub-phase review process may end. However, it is strongly recommended that the highlighted sections within external resources (extracted from the initial pool of thirty-one standards and guidelines) are used to better understand the interplay between the current sub-phase and its dependencies, low-level implementation details, etc. Without this, only a high-level understanding of sub-phase requirements is formed; this is insufficient to practically develop cyber incident response and recovery capability.

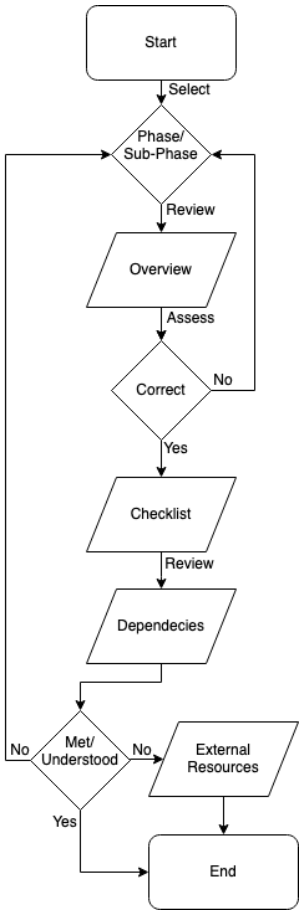

Fig. 4.3 Framework Process Flow

4.3 Framework Dependencies

During the stakeholder interviews discussed in Chapter 3, a common theme that emerged was the dependencies between different phases and sub-phases of the incident response life-cycle. It was noted that the success of several sub-phases within the mid-incident phase were dependent of the output of other phases. For example, during the stakeholder engagements, training of incident response teams was identified as a crucial step towards the success of mid-incident phases including Incident Detection, Reporting, Containment, Eradication and Recovery. To this end, the thirty-one standards and guidelines analysed in Chapter 3 were revisited to identify discussed dependencies between different phases and sub-phases of the incident response and recovery lifecycle. The results of this subsequent analysis have been compiled into Tables 4.1 and 4.2.

From this analysis, two major findings can be identified. Firstly, sub-phases from the Planning and Preparation phases of the framework indeed contribute to the success of sub-

	Planning					Preparation		
	RR	**RP**	**CA**	**TA**	**RM**	**Tra**	**RTA**	**ID**
Roles & Responsibilities					✓			
Response Planning	✓		✓	✓	✓		✓	
Criticality Assessment	✓			✓	✓			
Threat Analysis	✓				✓			
Risk Management	✓		✓	✓			✓	
Training	✓			✓	✓		✓	
Regular Testing & Auditing	✓				✓			
Incident Detection	✓	✓	✓	✓	✓	✓		
Resource Availability	✓	✓	✓	✓	✓			
Incident Reporting	✓	✓				✓		✓
Incident Containment	✓	✓	✓		✓	✓		✓
Incident Eradication	✓	✓	✓	✓	✓	✓		
Incident Recovery	✓	✓	✓		✓	✓		
Evidence Collection/Handling	✓	✓	✓		✓	✓		✓
External Relations Management	✓	✓			✓	✓		✓
Root Cause Analysis	✓	✓	✓			✓	✓	✓
Lessons Learnt	✓	✓						✓

Table 4.1 Response & Recovery Phase Dependencies (Part One)

phases within the Mid-Incident and Post-Incident phases and secondly, due to the cyclical nature of both the framework and the response and recovery lifecycle, Post-Incident sub-phases can contribute towards the success of Planning and Preparation sub-phases for future incidents. In particular, Roles and Responsibilities and Risk Management contribute towards the success of 100% and 81%, respectively, sub-phases within the framework; highlighting the importance of planning and preparation for successful incident response and recovery.

Risk Management, in particular, was identified as a crucial phase of the incident response and recovery lifecycle due to the outcome of this stage's effect on the way that subsequent phases are prioritised and implemented. Indeed, Risk Management corresponds to the identification, evaluation and prioritisation of risks including solutions for minimising, monitoring and controlling these. As part of this, assurance techniques are used to generate evidences as a means of making claims of assurance. While there includes many assurance techniques such as various review-types (documents, architecture, configuration, source code etc.), stakeholder interviews, formal verification, public review, and static and dynamic analysis, an increasingly popular method for evaluating and improving both cyber resilience and incident response capabilities is through adversary-centric security testing [93]. These engagements use highly specialised teams to emulate the actions of genuine malicious

	Mid-Incident							**Post-Incident**	
	RA	**IRep**	**IC**	**IE**	**IRec**	**ECH**	**ERM**	**RCA**	**LL**
Roles & Responsibilities									✓
Response Planning									✓
Criticality Assessment								✓	✓
Threat Analysis								✓	✓
Risk Management									✓
Training									✓
Regular Testing and Auditing									
Incident Detection									
Resource Availability									
Incident Reporting									
Incident Containment	✓	✓							
Incident Eradication	✓	✓	✓						
Incident Recovery	✓	✓	✓	✓					
Evidence Collection/Handling	✓	✓	✓						
External Relations Management	✓	✓	✓	✓	✓				
Root Cause Analysis		✓				✓			
Lessons Learnt	✓	✓	✓	✓	✓	✓	✓	✓	

Table 4.2 Response & Recovery Phase Dependencies (Part Two)

actors. Conducting such engagements helps organisations understand both the psychological factors and the techniques employed during genuine cyber attacks. In doing so, underlying vulnerabilities can be detected and patched, and incident response teams can be trained by being kept updated about tools and techniques used by modern attackers [148].

While adversary-centric security testing has become widely adopted within IT environments, this is not commonly the case for OT environments. Many training courses and certifications exist for IT penetration testing, such as the Offensive Security Certified Professional certification, whereas relatively little currently exists for OT penetration testing [210]. This presents several challenges, as not all tools and techniques used within IT environments apply to OT environments. For example, even simple actions such as active port scanning, often used during IT security tests, may result in system crashes within poorly configured OT environments [246]. As such, Chapter 5 aims to identify the current challenges of adversary-centric security testing within ICS/OT environments and areas for improvement in the context of cyber incident response and recovery preparedness.

Chapter 5

Current Challenges of ICS Adversary-Centric Security Testing

As discussed in Chapter 4, an essential aspect of an organisation's cyber security lifecycle, including defence and response, is preparation [264, 93, 259]. If organisations are not prepared to effectively defend and respond to cyber-attacks, whether targeting IT or OT, the impact of these can be disastrous and even possibly life-threatening in the case of most critical infrastructure environments [171]. Security testing, especially adversary-based such as red teaming, can provide significant benefits to ensuring that organisations are sufficiently prepared to defend and respond to cyber-attacks. Firstly, these types of engagement test current non-human-based defence and response capabilities by discovering vulnerabilities and weaknesses in existing protective measures such as firewalls and detection mechanisms. Secondly, they also test, train and improve human-based defence and response capabilities such as the incident response team or the general security culture of the organisation. While adversary emulation can prove to be more complex and costly than other engagements such as vulnerability scanning or training/exercising, since these engagements aim to be as close to reality as actual cyber attacks, this often results in a more thorough and in-depth understanding of current defence and response capabilities and ultimately leads to better improvement of these [148].

While conducting adversary-centric security tests has gained significant traction within IT environments, this is yet to be the case for OT [77]. The first reason for this is the only recent evolution of technologies within industrial environments. For example, the transformation of electrical grids into smart grids means that they now rely much more on both IT and ICS infrastructure than before [101]. Secondly, since OT systems are often found as part of underlying critical infrastructures, the critical nature of these environments means that teams

conducting any adversary-centric security test within these need to be highly specialised and vetted thoroughly, such as through the NCSC CHECK accreditation, for example [193].

This chapter extends the comparison made in Chapter 2 by detailing how the differences and similarities between IT and OT systems affect decision-making and actions taken during adversary-centric security tests within these environments.

5.1 Background and Related Work

Several surveys over the past decade, such as those by the US Government Accountability Office [276]or by Westby [281] on how board members and senior management within Critical Infrastructure govern the security of their organisation, have been conducted. Findings from these concluded that several security issues were missing at the time of the surveys, including an effective mechanism for sharing information on cyber security, general cyber security awareness, security features built into critical infrastructure networks, including OT, and metrics for measuring and assessing cyber security capabilities. While these surveys were conducted in 2011 and 2012, the studies conducted throughout Chapter 3 found similar results. This chapter concluded that while there have been significant advances in developing standards and guidelines for ICS and CNI, their widespread adoption was minimal, resulting in a less than complete picture. Preparation for incidents, including security assessments such as penetration testing, was identified as a crucial phase for effectively improving cyber incident response and recovery capabilities. Despite this, the use of adversary-centric assurance techniques was limited due to both the skill gap between OT engineering and general penetration testing and the limitations imposed by the safety-critical nature of ICS. A survey conducted by the SANS Institute showed that the top initiatives demonstrated by OT stakeholders included both performing security assessments or audits of control systems and their networks as well as initiatives to bridge the IT/OT gap [77]. While component testing, through assessment engagements such as vulnerability scanning, was considered a strong positive for improving network resilience, only 41% of participants claimed that they used these due to the risk of disrupting the operational process. A survey by Green et al. details the approaches adopted by security practitioners during risk assessment within ICS environments [93]. In this study, penetration tests were considered a distinct phase within the risk assessment process and could provide additional risk validation prior to appropriate mitigation. Scoping, including that of penetration tests, was identified as both one of the most important and most challenging parts of the risk assessment process.

Several works discuss and propose solutions for the general concerns raised from the presented surveys [42, 184, 259, 148, 93]. Conklin, for example, discusses the issues linked

with utilising IT-specific methodologies within an industrial context, especially concerning the CIA Triad [42]. The author proposes the addition of Resilience as an additional factor to consider alongside the CIA Triad while referencing several standards adapted from IT-specific security controls for use in OT environments such as NIST SP 800-53. While Conklin does not directly reference adversary-centric security testing within an industrial context, the fact that the CIA triad alone needs to be reconsidered when discussing OT environments demonstrates that further engagements to test this, such as penetration tests, also need to be reconsidered. Song et al. discuss the cyber risk assessment process for the design of I&C systems within nuclear power plants [259]. The final phase of their methodology recommends penetration testing to validate the proposed security design and implementation. However, the authors note that potential for disruption is possible when simulating attacks on the systems under consideration. No specifics are given on how these disruptions occur and what remediations exist for them. Murray et al. discuss the convergence of IT and OT and how this affects cyber security for critical infrastructures [184]. Using Hofstede's Theory, the authors demonstrate the differing cultural values between IT and OT across several dimensions, such as the Power Distance Index or the Uncertainty Avoidance Index. While the analysis does not provide insight into the technical differences between IT and OT, the cultural differences observed show that substantial readjustment is required to ensure the smooth transition to the convergence of the two technologies. Finally, Knowles et al. discuss assurance techniques for ICS, including penetration testing [148]. In this study, simulated security assessments are identified as being able to generate demonstrable audit evidence to assess and improve risk posture. Usage of security assessments was observed to be lacking due to the general absence of a workforce with specialised skills when assessing OT environments, especially those that are safety-critical.

Due to the high risk associated with causing additional overhead within an ICS/OT environment, such as through tools or techniques employed during active adversary-centric security testing, the majority of research conducted for security testing has been through moving the environment being tested away from the live environment [94, 83, 60] or development of specialised tools [11, 70, 232] for ICS/OT.

As a means of performing risk avoidance for testing of ICS, a majority of research on ICS security has focused on the construction of physical testbeds or digital twins. Green et al. propose a model for the design of ICS testbeds for this purpose [94]. Similarly, Gardiner et al. describe their lessons learnt from building an ICS and Industrial Internet of Things testbed [83]. The methodologies described in both of these papers provide a starting point for good practices when designing and developing ICS testbeds for security research. While these testbeds can be used to identify device-specific vulnerabilities and discover ICS-based

zero-days [95, 162], their generally lower-scale representation of live environments is better suited for host-level testing. It, therefore, makes it difficult to assess the full extent of an entire environment's security posture due to the many interactions between large groups of devices. However, one advantage of using ICS testbeds is that they can aid in determining the resilience of specific OT devices against tools and techniques that are planned to be used prior to an adversary-centric security test. This can be used to assess the risk these tools and techniques pose to a live environment without directly interacting with it. Similarly, digital twins, such as the one proposed by Dietz et al., for integration within SOCs [60] can also be used for similar purposes. However, while their virtual nature reduces the cost of development and increases the flexibility of implementation, they are generally less equipped for vulnerability research and instead used for simulations or direct monitoring.

Several specialised tools such as SimaticScan and PLCScan have been developed as part of an initiative to perform safe and efficient scans on ICS. PLCscan, for example, developed by Dmitry Efanov, is a tool written in python that is able to scan PLCs through Modbus or S7COMM [70]. This tool can query a range of data from the target PLC such as module name, firmware version, PLC name, serial number and more. However, no other functionality is possible; therefore, further assessment would need to be done manually or using other tools. Antrobus et al. identified the limitations of PLCScan and built upon it by proposing a Proof of Concept for SimaticScan [11]. The authors note that SimaticScan goes "beyond simply identifying potential vulnerabilities to verifying the existence of these vulnerabilities" for the target PLC. This is done through three distinct phases: reconnaissance scans, vulnerability assessment and fuzzing. The reconnaissance scan's functionality is similar to that of PLCScan in retrieving the PLC's information for CVE query alongside an SNMP scan. After this, SimaticScan is able to analyse PCAP files for identification of session IDs and plaintext vulnerabilities, perform a dictionary attack on any identified web server login forms, simulate a DoS attack on the PLC, simulate TCP hijacking, and verify unauthorised read/write access to the PLC Data Blocks. Finally, the tool can fuzz a PLC to determine other vulnerabilities. Overall, while the depth-of-testing of SimaticScan is extensive, its use is restricted to testing of Siemens devices only, severely limiting its effectiveness in environments that deploy devices from multiple vendors. As a means of aiding asset owners in selecting the appropriate tools for their environment and needs, Samanis et al. developed a taxonomy for contrasting ICS Asset Discovery Tools [232], which includes PLCScan. The taxonomy categorises the selected tools into three main classes: Specification, Execution and Output. Specifications of the tool detail its mode of operation, license scheme, scope, and supported protocols. The Execution category describes the tool's method of operation, its usage methodology, user interactivity, and approach to scanning. Finally, the Output category describes the

tool's output, such as listening ports, service identification, device info, deployment-specific information and vulnerability identification. Throughout this research, the authors note that none of these tools has information concerning their effect on the operational process; highlighting the need to perform a safety risk assessment prior to their utilisation. However, no methodology is provided for assessing each tool's risk to the operational process when being used as part of an adversary-centric security test.

To summarise, current research has mainly focused on reducing the risk of adversary-centric security testing by moving the testing environment away from the live environment towards testbeds and digital twins or developing specialised tools. However, there is little discussion on the extent of the risks that exist when performing tests within live environments. To this end, the following sections extend the comparison made between IT and OT in Chapter 2 to identify the challenges of conducting adversary-centric security tests within OT environments.

5.2 Analysis of Current Challenges

5.2.1 Methodology

As adversary-centric security testing aims to emulate actual cyber-attacks to test, train and improve an organisation's resilience, response, and recovery capabilities, the Tools, Techniques and Procedures (TTPs) used during these engagements closely mirror those of actual cyber-attacks. Although many of the cyber-attacks that have occurred over the years are somewhat unique, the adversaries behind them all follow, to some degree, the same steps for achieving their goals. Lockheed Martin mapped these steps to a framework titled the Cyber Kill Chain (CKC) [108]. As part of the Intelligence Driven Defense model, the framework identifies and details what adversaries must complete to ensure their objectives. The aim of this is for defenders to better understand the TTPs behind cyber-attacks to defend more effectively against them. The CKC is made up of seven steps; these are as follows:

1. Reconnaissance: Gain information on the target system by identifying and harvesting information that can be used to gain an initial foothold within the network.
2. Weaponisation: Create a payload to exploit the vulnerabilities found through reconnaissance.
3. Delivery: Deliver the payload to the target.
4. Exploitation: Gain access to the target by executing the payload to exploit vulnerabilities found through reconnaissance.

5. Installation: Establish a backdoor within the target to maintain access.
6. Command & Control: Open a command channel to be able to remotely manipulate the target system.
7. Actions on Objectives: Accomplish the attack's objectives.

It is worth noting that although the CKC contains phases similar to a linear process flow framework, it represents a dependency-based process flow. This means that the further an attacker advances through the kill chain, the more their subsequent actions depend on previously taken actions. Therefore, revisiting previous steps within the framework is extremely common and often essential, defining the CKC as more of a circular and non-linear process. For example, if an attacker has reached the Delivery stage (stage 3) of the CKC after crafting a payload to exploit a discovered vulnerability in the target network, they may need to conduct additional reconnaissance (stage 1) in order to discover how to deliver the payload as effectively as possible.

While the Lockheed Martin CKC provides a complete overview of the steps most adversaries take to conduct cyber-attacks, attacks on ICS require more depth and sophistication to succeed. Because of this, the SANS Institute developed the Industrial Control System Cyber Kill Chain [16]. This model, based on Lockheed Martin's original model, describes the steps taken by attackers to conduct a cyber-attack on ICS specifically. While simple ICS attacks such as industrial espionage or ICS disruption might not follow each stage of the ICS CKC, the steps described in this kill chain help defenders gain knowledge on how to better combat in-depth cyber-physical attacks, such as those originating from nation-state-sponsored groups. The ICS CKC is composed of two stages, each containing multiple phases; these are as follows:

- STAGE 1: Cyber Intrusion Preparation and Execution
    1. Planning: Reconnaissance.
    2. Preparation: Weaponisation and Targeting.
    3. Cyber Intrusion: Delivery, Exploitation, and Installation/Modification.
    4. Management and Enablement: Command & Control.
    5. Sustainment, Entrenchment, Development and Execution.
- STAGE 2: ICS Attack Development and Execution
    1. Attack Development and Tuning.

2. Validation and Testing.
3. ICS Attack: Deliver, Install/Modify, and Execute.

As we can see, the first stage of the ICS CKC closely resembles the Lockheed Martin CKC. The two models start to differ after this, however. The ICS CKC contains an additional stage because successful attacks on ICS with re-attack options require extremely high levels of confidence to execute.

Although both the Lockheed Martin CKC and the ICS CKC provide a holistic overview of the steps used during an adversary-centric security test, due to the high-level nature of these models, they provide little technical detail on the TTPs used during each phase of the CKC. To provide more technical depth to the analysis, both the MITRE ATT&CK and MITRE ICS ATT&CK Frameworks were, therefore, selected for this [176, 178]. These frameworks provide a knowledge base of adversary tactics and techniques based on real-world observations. Each TTP is categorised by attack types such as reconnaissance or lateral movement. These frameworks aim to provide defenders with knowledge on the TTPs used by attackers to understand them better and, consequently, better defend against them. As demonstrated in Chapter 2, there are distinct differences that need to be considered when attacking ICS networks compared to traditional IT networks. The following sections will detail these differences when conducting an adversary-centric security test on IT and OT systems. While the phases of the CKC and TTPs detailed within the ATT&CK frameworks are often leveraged during an adversary-centric security test to emulate real-world adversaries, a specific subset of these TTPs and phases may be utilised depending on the type of engagement being done. For example, a red team engagement is likely to leverage all phases of the CKC, while a typical vulnerability scan may make use of the reconnaissance phase and part of the weaponisation phase only.

5.2.2 Results

Reconnaissance

At the beginning of any adversary-centric security test, reconnaissance must be conducted to gain the information required to exploit the target systems. Two types of reconnaissance exist: passive reconnaissance and active reconnaissance. Passive reconnaissance refers to conducting reconnaissance that does not directly interact with the target system. This can either correspond to using non-technical reconnaissance such as Open-Source Intelligence (OSINT) through search engine searches (google, Shodan) and tools such as Netcraft or using passive tools and methodologies such as network sniffing. Active reconnaissance,

on the contrary, directly interacts with the target system to obtain information on it. This corresponds to using techniques and tools such as vulnerability scanners, port and service scanners, fingerprinting, banner grabbing, etc. While reconnaissance is integral to any cyber intrusion exercise, only the MITRE ATT&CK Framework provides specific TTPs for this phase. This is because the ICS ATT&CK framework is defined as a "knowledge base [that] can be used to better characterise and describe post-compromise adversary behaviour" rather than a framework encompassing the whole CKC. Despite the TTPs described in the ATT&CK framework being meant for IT networks, most of them can also be applicable for OT networks. The MITRE ATT&CK framework details four different types of goals when conducting general reconnaissance: Gathering victim host (T1592), identity (T1589), network (T1590) and organisation (T1591) information. Once enough actionable information is acquired on these, teams can proceed to the next step of the CKC, the weaponisation of a payload. The following sections discuss the characteristics of passive and active reconnaissance techniques and TTPs, and how they differ between IT and OT environments.

Passive Reconnaissance

As mentioned, passive reconnaissance is a means of acquiring actionable information on a target system or network without directly interacting with it. While it often takes considerably longer to obtain valuable information through this method, the fact that no direct interaction with the target is required means that detection is infrequent. This presents several opportunities for attackers and red teams alike, including the freedom of evading detection, gaining considerable time to develop exploits and more. While the advantages of conducting passive reconnaissance are plentiful, this information can be somewhat limited depending on the context. IT-based targets often have a plethora of public-facing information available for attackers to use through tools such as email harvesters, domain lookup, search engine dorking and more. The MITRE ATT&CK framework details several TTPs associated with conducting passive reconnaissance. This includes the use of searching through closed sources (T1597), such as searching through or purchasing private data, including technical data from threat intelligence vendors and other private sources; and searching through open sources such as technical databases (T1596), open websites and domains (T1593), and victim-owned websites (T1594).

While it is also possible to use these methods to conduct passive reconnaissance on OT-based networks, these often mean that much of the information that is of value to an attacker is hidden from the public domain. In recent years, many ICS networks have started integrating IoT to improve automation, data collection and more. Despite the benefits this provides, this has also significantly increased the potential attack surface for these by, in

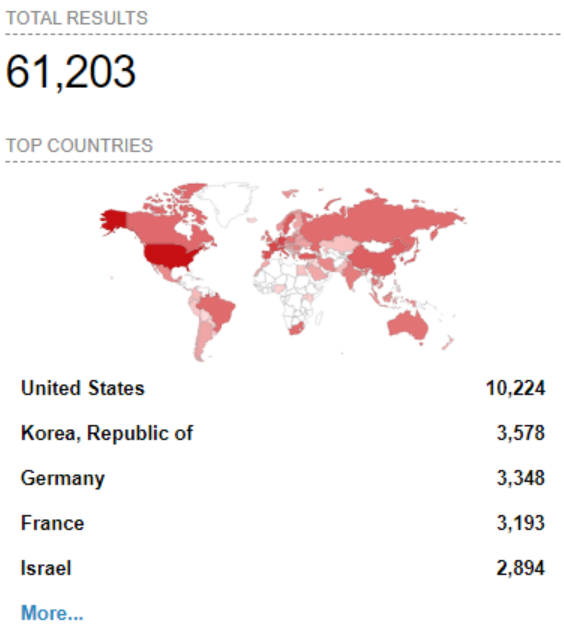

Fig. 5.1 Shodan "port:502" Search Results

several cases, making these networks public-facing. For example, a simple search using the Shodan search engine shows that over 61 000 public-facing devices are running with port 502 open, which commonly uses the Modbus TCP/IP protocol [165]. Figure 5.1 illustrates the results of using Shodan to search for devices with port 502 open. If operators do not correctly setup their Modbus TCP/IP connections, this can be exploited by attackers with relative ease due to the large number of vulnerabilities that exist within the default version of Modbus TCP/IP, including the use of clear text, the lack of integrity checks, and the lack of authentication [18].

Overall, while the amount of actionable intelligence gained from performing passive reconnaissance can vary from organisation to organisation, the primary objective is to gain information while preventing detection. As such, engagements that include the participation of a red team can benefit greatly from this. However, publicly available information on industrial networks, such as those within critical infrastructures, is often and should be relatively limited compared to what could be gained from an IT-based organisation. These findings are summarised in Table 5.1.

Active Reconnaissance

As opposed to passive reconnaissance, active reconnaissance is used to obtain information on a target system through direct interaction. The most common way of conducting this

	Information Technology	**Operational Technology**
Detectability	Low	Low
Disruptability	Low	Low
Potential Gathered Information	High	Low
Ease of Information Gathering	High	Low-to-High

Table 5.1 Passive Reconnaissance Summary

type of reconnaissance is through port and vulnerability scanning (T1595) using manual tools such as Nmap or automated scanners such as Nessus [161, 268]. While such tools can help attackers and red teams obtain valuable information for discovering vulnerabilities and weaknesses within target systems, misusing them can lead to an extremely high risk of detection. This includes, for port scanners, using TCP connect scans, aggressive scan timings, and aggressive scripts, which are all often detected by IDSs and IPSs. It is, however, possible to configure scans to minimise detection risk using techniques such as packet fragmentation, decoy scanning, source IP address spoofing, source port spoofing, and lowering scan timing. A difference between IT and OT systems is the risk that active reconnaissance can have on OT systems, especially if they are legacy devices. Due to how these systems are programmed, any form of unrecognised or heavy network traffic could cause software crashes, resulting in significant operational impact; as demonstrated when a ping sweep on an active ICS network to identify all hosts caused a fabrication plant to hang, destroying $50 000 worth of equipment [64]. As such, additional care must be taken when performing active scanning on OT systems, such as reducing scanning speed to minimise network traffic, performing the correct scan type for the target system (i.e. avoiding UDP scans on devices using TCP ports), and manually selecting scripts to run. Many standards and guidelines, such as the IEC 62443 series, recommend including subject specialists such as OT engineers when performing any cyber security activity, including adversary-centric security testing [119, 115]. In recent years, custom scanners, like PLCSCAN and SimaticScan, have been designed to facilitate performing active reconnaissance on devices such as PLCs [70, 11]. While these allow for better and easier scanning of specific ICS types such as Simatic PLCs, these solutions are not comprehensive. Because of this, teams conducting reconnaissance will often resort to primarily using passive reconnaissance or more generalised tools like the ones discussed here.

Another form of active reconnaissance is using social engineering (T1598) to obtain information such as credentials or private information on networks or systems. Similar to active scanning, properly conducting social engineering is crucial for evading detection. Any indication that a third-party message, such as an email, is being used as social engineering can

	Information Technology	Operational Technology
Detectability	High	High
Disruptability	Low	High
Potential Gathered Information	High	High
Ease of Information Gathering	High	Low

Table 5.2 Active Reconnaissance Summary

be detected by both automated methods and potential victims. Messages need to, therefore, be carefully crafted to avoid such detection, often done using social-engineering tools such as the TrustedSec Social-Engineer Toolkit [142].

Overall, while active reconnaissance can obtain a large amount of actionable information in little time, if not used properly, it can easily be detected and, in some cases, even cause accidental operational impact to systems. To prevent this, time needs to be taken to fully understand the tools being used, and a near-comprehensive understanding of the targeted devices is required, making black-box testing considerably more difficult, albeit not impossible. These findings are summarised in Table 5.2.

Weaponisation

Once enough actionable information has been acquired from reconnaissance, the weaponisation of a payload to exploit the target system can begin. While it is entirely possible to craft a payload manually, attackers often make use of automated tools such as the Metasploit framework [226] or vulnerability databases such as NIST NVD [203] or MITRE CVE [175] to then obtain Proof of Concept code to modify. During this stage, attackers will often combine a Remote Access Trojan (RAT) for Command & Control post-exploitation and the malicious code used to exploit vulnerabilities discovered during reconnaissance into a deliverable such as client data application files (PDF, Word, etc.).

The weaponisation of a payload often depends on the attacker's primary goal, which is often associated with the desired impact of the attack. Both the MITRE ATT&CK and ICS ATT&CK Frameworks categorise four main impact types of cyber-attacks: Denial, Sabotage, Collection, and Control. For adversary-centric security testing, distinguishing how different exploits can impact systems is essential for calculating risk and identifying relevant mitigation strategies.

Denial is often the most common type of attack impact due to its ease of execution. These attacks aim to deny the target of either view, access, or control to their own environment. Such attacks that cause denial include DoS attacks (T1499, T1498), encryption attacks (T1486),

account access removal (T1531), service stops (T1489) and system shutdowns (T1529). While these attacks cause operational downtime, they often only have a short-term impact on IT systems since, in most cases, there is no destruction of data or hardware. However, for OT systems, even slight downtime can be detrimental for time-sensitive environments such as power plants.

While a consequence of sabotage can also be denial, the act of sabotage involves the deliberate destruction or obstruction of regular operation. Such attacks that cause sabotage include the destruction of data (T1485) through wiping disks (T1561), for example; corrupting firmware (T1495); damage to property (T0879); and, in extreme cases, a loss of protection or safety (T0837, T0880). Similar to denial, the consequences of attacks causing sabotage are vastly different between IT and OT systems. Sabotage attacks can have severe economic and social consequences for IT systems if proper recovery steps are not implemented. For example, the corruption of a database containing customer data could lead to a complete halt in operations, leading to potential severe economic loss and the potential loss of existing customers due to dissatisfaction or distrust. While sabotage attacks on OT systems can also have economic and social consequences on future operations, the potential for loss of protection or safety makes it critical to effectively respond and recover from such attacks, as such consequences could lead to a danger to life.

Any attack involving collection corresponds to actions where information or data is extracted from the target system. The goal of such an attack can include further reconnaissance (see section 5.2.2), the theft of data to sell on black markets or use for blackmail. The ATT&CK Framework categories 17 different collection techniques, including capture-based techniques such as screen (T1113), video (T1125), audio (T1123), clipboard (T1115), and input capture (T1056); direct data extraction techniques from various sources such as cloud storage (T1530), configuration repositories (T1602), information repositories (T1213), local systems (T1005), network shared drives (T1039), removable media (T1025), and mail servers (T1114). While collection attacks do not have a direct operational impact on either IT or OT systems, it leaves the target organisation open to further action from adversaries. These can be especially devastating for national critical infrastructures when they are the victim of industrial espionage from other nation-states, for example.

In most cases, control attacks are considered the most dangerous attack types to affect target systems for both IT and OT. For these attacks, adversaries gain remote code execution (RCE) on their target. While crafting the payload to gain RCE on a target is done during the weaponisation phase, Command & Control is detailed in a separate phase of the CKC and is therefore discussed more in-depth in section 5.2.2.

	Information Technology	**Operational Technology**
Denial	Loss of Availability/Revenue	Loss of Availability/Control/Safety
Sabotage	Loss of Revenue/Data	Loss of Data/Safety
Collection	Theft of Data	Theft of Operational Information
Control	Manipulation of Control	Manipulation of Control/Loss of Safety

Table 5.3 Impact differences between IT and OT systems

Table 5.3 summarises the different impact types and their consequences on IT and OT, respectively.

Delivery

Once a suitable payload has been created for exploiting the target system, it needs to be delivered to the victim so that it can be executed. Two main techniques exist for doing this: adversary-controlled delivery and adversary-released delivery.

Delivery of a payload classified as adversary-controlled corresponds to a payload that executes through direct execution of an adversary. This is often done when remote access to the target is possible through open ports. For example, adversaries may access a system by exploiting public-facing applications (T1190) such as websites, databases, and standard services. Physical delivery of payloads is also possible using replication through removable media (T1091), such as by taking advantage of the autorun feature of most devices when inserting a USB drive. While this method may seem less viable due to strict physical security measures implemented within critical infrastructures such as power facilities, it is still possible through a trusted user, for example, as demonstrated in the Stuxnet attack of 2010 [199].

When attackers cannot directly access the target system, they may resort to delivering their payload through adversary-released means. When using this technique, adversaries often use social engineering to trick unsuspecting users into executing a payload. This can either be done through drive-by-compromise (T1189) by compromising a website that a user visits throughout normal browsing, for example; through direct phishing tactics (T1566) such as providing malicious links or attachments in emails or messages; or even through supply chain compromise (T1195) by inserting malicious code into tools used by the target organisation. Supply chain compromise, in particular, has gained much attention recently due to the 2020 global supply chain cyber attack, which affected around 18,000 different organisations using software distributed by SolarWinds, including the United States National Nuclear Security Administration [288, 137, 54]. Such an attack demonstrates that even

Technique	ATT&CK Technique
Adversary-Controlled Delivery	Public-Facing Applications (T1190), Replication through Removable Media (T1091)
Adversary-Released Delivery	Drive-by-Compromise (T1189), Phishing (T1566), Supply Chain Compromise (T1195)

Table 5.4 Summary of Delivery Techniques

though the initial attack targeted IT systems, a wide variety of critical infrastructures, with some being comprised of OT, were affected.

Despite both IT and OT environments being vastly different in function and architecture, the techniques used to deliver malicious payloads for execution are often similar, as demonstrated by the listed techniques under "Initial Access" in both the MITRE ATT&CK and ATT&CK ICS frameworks. Some techniques do differ for OT-specific systems such as Data Historian Compromise (T0810) in the case of the ICS framework; however, general techniques such as the ones discussed here (drive-by-compromise, phishing, etc.) and summarised in Table 5.4 apply to both IT and OT environments and often require little to no modification in terms of methodology.

Exploitation

Once the malicious payload has been successfully delivered onto the target network or system, execution of the code to exploit the target can begin. This phase of the CKC uses the weaponised payload discussed in Section 5.2.2. While the weaponisation stage of the CKC is difficult to detect and mitigate due to the activities during that stage being entirely separate from the target systems, if a malicious actor has reached the exploitation stage of the CKC, this should be reported as an incident, and appropriate response and recovery actions need to be taken, including during a red team engagement.

While, ultimately, the goal of conducting response and recovery is to return to a state of normal operation, different methods to achieve this outcome are required depending on the environment. For IT systems, the conservation of the CIA Triad is considered a high priority when responding to a cyber incident [79]. Confidentiality is vital for IT-based organisations to recover, as the unauthorised sharing of private data can have severe economic consequences and damage public relations. Violating the General Data Protection Regulation (GDPR), for example, can bring about severe fines of up to 20 million euro or 4% of worldwide turnover for the preceding financial year regardless of cause, including data breaches caused by cyber attacks [269]. This was the case for British Airways, which had to pay a fine of

20 million pounds sterling in 2020 due to a data breach in 2018 where malicious actors obtained private information such as log-in details, payment card information, and customer addresses [273]. Integrity also plays an important part when responding to IT cyber incidents, as data tampering is likely during such an event. Recent findings have reported that cyber attacks with the sole intention of manipulating data have increased significantly since 2020, leading to the spread of disinformation [253]. Compromising data integrity can also serve as a method of detection and defence evasion through techniques such as manipulating indicators (T1070) which can prevent defenders from properly using event collection and reporting. Similarly, data is often rendered useless without availability as it cannot be shared with intended users. Attacks that cause denial or sabotage, as discussed in Section 5.2.2, can affect the availability of systems.

While the CIA triad is considered a staple model for developing IT security policies and, consequently, something that should be tested thoroughly when conducting adversary-based tests, this is not always the case for OT-based systems. Due to the time-critical nature of these environments, availability is allocated considerably more priority than confidentiality and integrity (however, availability can be dependent on integrity). Furthermore, due to the operational nature of these environments, safety considerations also play a critical part when testing for security resilience. For example, during an attack on a German steel mill in 2014, adversaries gained access to a blast furnace control mechanism, preventing it from shutting down and causing significant damage to the machine itself and the surrounding environment [154, 27]. While there were no human casualties, a loss of safety was observed. Therefore, these differences, summarised in Table 5.5, need to be considered when conducting an adversary-centric security test to identify appropriate risk mitigation techniques. This table provides relative priority for each category of the CIA triad. This signifies that, for IT, while Availability could be considered a high priority for specific applications such as streaming services, the financial impact of having confidentiality compromised is still considered higher than if availability were to be compromised. Therefore, in general, confidentiality is allocated higher priority than availability within IT environments. This same reasoning is applied for the prioritisation of the CIA triad within OT environments.

Installation, Command and Control, and Actions on Objectives

Once a discovered vulnerability has been exploited, an adversary, depending on their goal, will then seek to install further capabilities such as persistent remote access or an escalation of privileges. This will often involve revisiting previous steps, such as reconnaissance, to obtain further information on how this can be done. A phenomenon often observed with attacks targeting ICS includes malicious actors gaining access to the industrial zone by pivoting

CIA Triad Category	Priority for IT	Priority for OT
Confidentiality	High	Low (Medium/High for manufacturing processes that include corporate secrets such as chemical recipes)
Integrity	Medium	High (Due to effect of Integrity on Availability)
Availability	Low	High (Due to potential in reduction in Safety)

Table 5.5 CIA Triad Prioritisation Summary

from the enterprise zone. While this phase of the CKC is often not required for traditional penetration testing, demonstrating further capabilities provides additional depth for advanced security testing such as red team engagements.

To this day, achieving RCE on ICS has only been observed in highly advanced attacks [171]. For this reason, minimal detail is given on the TTPs provided by the MITRE ICS ATT&CK Framework. The framework details three techniques for this which are Commonly Used Ports (T0885), Connection Proxy (T0884), and Standard Application Layer Protocol (T0869). In contrast, the enterprise framework details 16 different categories of TTPs used in Command and Control activities, illustrating that knowledge of Command and Control in ICS is still relatively limited to this day. Recent research on Process Comprehension at a Distance has demonstrated the possibility of RCE by leveraging unused memory within PLCs, effectively creating a covert Command and Control channel for realising further actions on objectives [95]. However, more traditional methods for remotely controlling OT exist due to legacy design decisions in PLCs. For example, simply gaining network access could lead to RCE by leveraging the poor use of access control within standard industrial protocols and directly interfacing with PLCs.

Stage 2: Development and Execution (ICS only)

Due to the high confidence required for conducting precise cyber attacks and thorough adversary-centric security tests on ICS, the ICS CKC contains an additional stage to the traditional CKC. During this stage, attackers use the knowledge they gained from the previous stage to develop and test their capabilities so that a high-confidence attack on ICS can be carried out. If an impact is observed during the first stage of the ICS CKC, this is unintended and often caused by equipment failing due to its sensitivity. During this phase, several TTPs can be used to further increase the impact and precision of an attack by inhibiting response functions or impairing process control. For example, attackers may block or spoof

reporting messages (T0804 & T0856) to delay response and recovery actions. Because of the precise requirements for this stage, the overall timeline for developing an attack of this capability is often greatly extended compared to low-confidence or imprecise attacks. This development time requirement can be directly translated to the time required for conducting an adversary-centric security test, often being a red team engagement at this stage based on scoping constraints such as cost and time.

Security Testing Software and Tools

As demonstrated throughout this Chapter, the differences between IT and OT fundamentally affect how adversary-centric security tests are performed within these environments. Because of this, the tools and techniques employed throughout the different phases of the security testing life-cycle also need to be considered. For example, during reconnaissance (discussed in section 5.2.2) blind scanning in an IT environment may be used for discovering running services on devices, but doing so within an industrial environment could lead to disruption due to unknown protocol compatibility issues. Because of this, specialised tools for security testing within OT environments have been developed. For example, the ControlThings Platform is a specialised penetration testing distribution for ICS [241]; similarly, Kali Linux is a penetration testing distribution for traditional IT [242].

Due to the often proprietary nature of software and protocols used within OT environments, security testing tools for OT are designed for the testing of specific protocols or products. For example, PLCScan is a tool developed by Dmitry Efanov for retrieving information on PLCs that use Modbus or S7comm [70]. Similarly, SimaticScan, developed by Antrobus et al., can only scan Siemens-based PLCs but offers more depth of testing by also being able to scan for known vulnerabilities and perform fuzzing to discover unknown vulnerabilities [11]. To contrast the differences in functionality of specialised tools for ICS security testing, Samanis et al. developed a taxonomy for categorising ICS Asset Discovery Tools [232]. This covers security testing tools for asset discovery only, and the significant disparity in functionality and practicality between these tools demonstrates the challenges in the development and useability of software used for security testing within ICS environments.

Overall, the most noticeable difference between IT and OT security testing tools/software is that those developed for enterprise security testing often offer extensive functionality for specific tasks, whereas OT-specific tools also need to consider the compatibility of non-standard software or protocols, which often limits their applicability. To summarise, Table 5.6 presents an example list of tools used for security testing in IT and OT environments and comments on the challenges involved.

Functionality	IT Tools	OT Tools	Comments
Security Testing OSs	Kali Linux, ParrotOS	ControlThings	While traditional security testing distributions can be used for OT environments, a thorough understanding of the effect of tools available from these is required to prevent disruption
Port Scanning	Nmap, Netcat, Zenmap	Nmap, Netcat, Zenmap	Port scanning often unsuitable for OT environments due to potential compatibility issues
Passive Network Enumeration	Wireshark, TCPDump	Wireshark, TCPDump, NSA GRASSMARLIN	While less precise as active scanning, passive enumeration is preferred for OT due to low risk of disruption
Vulnerability Scanning	Nessus, OpenVAS	PLCScan, SimaticScan	OT vulnerability scanners are often very limited in what devices they can be used on
Exploitation Frameworks	Metasploit, CORE IMPACT, Immunity CANVAS	Industrial Exploitation Framework, ICSSPLOIT	Custom exploitation in OT environments often favoured to increase precision and stealthiness of attack
C2 Frameworks	Empire, Covenant	Custom Capabilities	C2 still in infancy for OT but is possible as shown by recent research [95]
Adversary Emulation Frameworks	Cobalt Strike, CALDERA	OT CALDERA	Adversary Emulation Frameworks for OT still in development and not open source due security concerns

Table 5.6 Example Software/Tools used for Security Testing IT and OT

5.3 Testbed Experimentation

5.3.1 Methodology

Throughout this Chapter and Chapter 2, the observed differences between IT and OT demonstrated that these need to be carefully considered prior to conducting any active form of adversary-centric security test within industrial networks for improving R&R capabilities. To evaluate these findings, several experiments have been conducted on the Lancaster University ICS testbed, which was previously used for the development of synthetic attack scenarios as part of the stakeholder engagement in Chapter 3, Section 3.3. As a reminder, the testbed has been built using physical, real-world hardware and software produced by major ICS vendors, including Siemens, Schneider, Allen Bradley, and ABB and is actively being used to support the development and evaluation of industry driven tools [92, 94]. This, therefore, provides a high degree of realism for experimentation.

To identify to what extent active penetration testing techniques affect OT operations, several tools with varying degrees of risk in terms of affecting availability were selected. While these techniques may not necessarily be used for all types of adversary-centric security tests, they were selected based on their potential to disrupt operational processes within an industrial network by affecting network traffic or endpoint resource usage. The techniques used in the experiment are as follows:

- Default Ping Sweep: control test.
- Ping Flood: medium network traffic test.
- Hping3 Flood: heavy network traffic test.
- Malformed Packet Ping: abnormally large packet size test.
- Low-Risk Nmap Scan: TCP connect scan with 1 second delay between probes on top 1000 ports.
- Medium-Risk Nmap Scan: connect scan on all TCP ports, default speed and no probe parallelisation.
- High-Risk Nmap Scan: scan on all TCP/UDP ports, fastest speed, OS detection, version detection, script scanning, and traceroute.
- Nessus Scan: commonly used vulnerability scanner test.

All of the selected techniques are primarily used during the reconnaissance phase of an adversary-centric security test. Although techniques used during subsequent phases of the CKC, such as exploitation, can also adversely affect the operational process. As described in Sections 5.1 and 5.2, the tools used for this depend greatly on identified vulnerabilities that can be unique to each device. Therefore, using reconnaissance techniques and their subsequent tools provides consistency when testing on distinct targets.

Four devices were selected for experimentation to identify how the usage of these techniques could affect operational processes within an industrial environment. Firstly, to test legacy OT, a Siemens SIMATIC ET-200S was selected, which is, to this day, still commonly used in industry [250]. These PLCs started production in 1994 and, as of the 1st of October 2020, are currently in product phase-out with a total phase-out planned for 2023. Next, to test more recent PLC lines, the Siemens SIMATIC S7-1200 was selected. Initially released for delivery in 2009, the S7-1200 currently has no announced phase-out date and has improved system properties over the older S7-300 and 400 series PLCs to meet the requirements of modern OT environments. To understand the effect of the selected tools on different PLC brands, an Allen-Bradley Logix5561 was also selected for the experiment. Finally, to demonstrate how these techniques could affect OT devices compared to IT devices, the tools and techniques were also tested on an IT workstation used to modify and upload code to the PLCs within the testbed.

Due to the significant differences in uses between the selected OT and IT devices, their technical specifications conform to the requirements of their end-users and are therefore described in vastly different terms. For example, the product details of the selected PLCs focus more on environmental resilience such as interference immunity, maximum air pressure operation, relative humidity operation etc., as opposed to the traditional and more IT-focused description of the capabilities of a device's components such as power usage, CPU clock rate, RAM clock speed etc. To this end, limited information can be inferred when directly comparing hardware specifications between IT and OT. This is shown in Table 5.7 which provides technical specifications for each of the selected devices used in the experiment based on data sheets provided by their respective vendors and internal system information. Not only is the terminology between IT and OT vastly different, but cross-OT vendor terminology also presents significant challenges for conducting a direct comparison. For example, while the work memory in SIMATIC PLCs can be defined as equivalent to RAM for an IT Workstation, the CPU details of each of the selected devices make it difficult to make a quantitative comparison between their respective speed and efficiency; further amplifying the IT and OT gap demonstrated in Chapter 2, and this Chapter.

	ET200S	**S7-1200**	**AB Logix5561**	**IT Workstation**
Power/Current Draw	320mA @ 24V DC	1.2A @ 24V DC	14mA @ 24V DC	290W
Memory	128KB (work) + optional load	50KB (work) + 1MB (load)	478KB (I/O) + 2MB (user)	16.0GB (RAM)
CPU Speed	3μs/instruction (float)	2.5μs/instruction (float)	100 programs/task (32 concurrent max)	1 core @ 3.30GHz
OS/Firmware	IM151-8 PN/DP V2.7.1	1212C V3.0.2	1756-L61S V10.007	Windows 7 Enterprise V6.1.7601

Table 5.7 Device Hardware Specifications

To evaluate how these techniques could adversely affect availability, two metrics were selected, each split into sub-metrics:

- Network Delay:
  - Maximum Round-Trip Time - the maximum possible effect on availability.
  - Average Round-Trip Time - the average effect on availability.
  - Packet Loss - the amount of total availability loss.
- CPU Resource Usage:
  - Maximum CPU Job Execution Time or Usage - the maximum load increase on the CPU.
  - Average CPU Job Execution Time or Usage - the average load increase on the CPU.
  - CPU Response - the response rate of the CPU.

A default ping scan was conducted in parallel with running the selected tools to collect data on network delay. For collecting data on the CPU usage of each tested endpoint, custom python (for the PLCs) and Powershell (for the IT workstation) scripts leveraging the protocols used by these (S7comm, HTTP, Ethernet/IP) were utilised.

5.3.2 Results

The results from running the selected tools on the four targets can be found in tables 5.8, 5.9, 5.10 and 5.11. During the entirety of the test, the ET-200S' availability was greatly affected by the more aggressive tools such as the hping3 flood, the high-risk Nmap scan, and the Nessus scan; resulting in a near-total loss of availability through resource overload in the case of the hping3 test or full system crashes for both the high-risk Nmap scan and the Nessus scan. All three of these techniques generated significant network traffic, resulting in the PLC being unable to reply to these on time. During the Nmap and Nessus scan, vulnerability and network enumeration scripts were performed, resulting in the PLC crashing due to its

	Max RTT	Avg RTT	Packet Loss	Max CPU Time	Avg CPU Time	CPU Response
Default Ping	13.203ms	5.133ms	0%	27ms	18.13ms	100%
Ping Flood	20.382ms	8.898ms	0%	38ms	25.41ms	100%
Hping3 Flood	1397.1ms	424.140ms	82.6087%	N/A	N/A	0%
Malformed Ping	7.444ms	4.629ms	0%	28ms	19.38ms	100%
Low-Risk nmap	11.276ms	4.331ms	0%	32ms	19.12ms	100%
Medium-Risk nmap	11.617ms	4.345ms	0%	49ms	27.74ms	100%
High-Risk nmap	227.995ms	33.658ms	CRASH	813ms	90.83ms	CRASH
Nessus Scan	283.754	49.950ms	CRASH	919ms	46.94ms	CRASH

Table 5.8 SIMATIC ET-200S Experiment Results

	Max RTT	Avg RTT	Packet Loss	Max CPU Time	Avg CPU Time	CPU Response
Default Ping	2.018ms	0.958ms	0%	13ms	11.04ms	100%
Ping Flood	1.586ms	0.680ms	0%	22ms	19.36ms	100%
Hping3 Flood	1.463ms	1.088ms	95.45%	14ms	12.67	26.32%
Malformed Ping	2.216ms	0.933ms	0%	14ms	11.32ms	100%
Low-Risk nmap	2.312ms	1.056ms	0%	14ms	11.13ms	100%
Medium-Risk nmap	2.052ms	0.643ms	0%	14ms	10.98ms	100%
High-Risk nmap	2.326ms	0.846ms	0%	20ms	11.68ms	98.86%
Nessus Scan	2.649ms	0.766ms	0%	28ms	12.15ms	100%

Table 5.9 SIMATIC S7-1200 Experiment Results

	Max RTT	Avg RTT	Packet Loss	Max CPU Load	Avg CPU Load	CPU Response
Default Ping	1.790ms	0.750ms	0%	1.1%	0.81%	100%
Ping Flood	0.558ms	0.381ms	93.18%	1.8%	1.55%	100%
Hping3 Flood	823.89ms	813.78ms	98.53%	N/A	N/A	0%
Malformed Ping	1.797ms	0.786ms	0%	1.1%	0.81%	100%
Low-Risk nmap	1.792ms	0.714ms	0%	1.2%	0.92%	100%
Medium-Risk nmap	0.688ms	0.440ms	0%	7.9%	5.65%	100%
High-Risk nmap	1.798ms	0.653ms	0%	12.7%	3.25%	100%
Nessus Scan	0.676ms	0.418ms	0%	31.1%	1.11%	100%

Table 5.10 Allen-Bradley Logix5561 Experiment Results

	Max RTT	Avg RTT	Packet Loss	Avg CPU Load	CPU Response
Default Ping	1.212ms	0.735ms	0%	18.8%	100%
Ping Flood	0.809ms	0.501ms	0%	24.85%	100%
Hping3 Flood	5.442ms	2.585ms	8.08%	40%	100%
Malformed Ping	1.172ms	0.747ms	0%	18.68%	100%
Low-Risk nmap	1.189ms	0.715ms	0%	18.71%	100%
Medium-Risk nmap	1.573ms	0.707ms	0%	14.38%	100%
High-Risk nmap	1.226ms	0.664ms	0%	14.78%	100%
Nessus Scan	3.658ms	0.684ms	0%	14.23%	100%

Table 5.11 Windows 7 Workstation Experiment Results

inability to handle these correctly. While the ping flood did not result in a total loss of availability, it caused the network delay and the CPU response time to increase. Based on an organisation's requirements for availability, including system dependencies, this can cause adverse consequences on the overall network. For example, in a time-sensitive environment such as the safety systems within a nuclear power plant, this decrease in availability would be undesirable. However, in less time-critical environments, tools with similar throughput may be acceptable for use. While no increase in network latency was observed for the medium-risk Nmap scan, the CPU response time did increase by 50%, which, similar to the high-risk Nmap scan, could be undesirable depending on the environment. No significant increase in network delay nor CPU response time was observed for the remaining tests.

Due to the more-recent hardware and firmware in the S7-1200, less impact on its availability was observed during the experiment than the ET-200S. A significant decrease in availability was observed when performing an hping3 flood. However, either a negligible decrease or no change in availability was observed for the seven other tests, including the Nessus and the high-risk Nmap tests. While these two tests did not increase network latency, an increase from 13ms to 20ms and 28ms respectively in CPU response time was observed, which may be undesirable for specific environments. This demonstrates that, apart from the most aggressive techniques, most tools generally present less risk to the availability of the S7-1200 than the ET-200S.

Despite the Logix5561 having more hardware resources than the S7-1200 and the ET-200S, it performed considerably worse than these when running heavy network generating tools such as ping and hping3. Both the ping flood and the hping3 flood saw a near-total loss of availability as opposed to the S7-1200 and ET-200S, which only saw a considerable loss of availability when running the hping3 flood. Despite this, the six other tests, including the high-risk Nmap scan and the Nessus scan, resulted in a negligible increase in network delay. However, these two tests, in particular, did produce a noticeable increase in the CPU usage (12.7% and 31.1%, respectively) for the Logix5561.

Findings from running the tools on the IT workstation show that it is generally more resilient than the tested PLCs. Only the hping3 flood resulted in a slight decrease in availability, although no total loss was observed. However, this is expected as an hping3 flood can generate over 170,000 packets per second which could affect even the most resilient of systems without proper DoS mitigation techniques. All seven other tests had a negligible effect on the workstation.

5.3.3 Discussion

Overall, the results from the experiment demonstrate that adversary-centric security testing techniques generally affect the availability of OT equipment more than IT. However, the extent to which these techniques affect OT is less than described by the findings from theory. This, therefore, signifies that adversary-centric security testing is indeed possible within OT environments, following proper risk quantification of the effect on availability that techniques used have on the systems under consideration. Furthermore, while legacy OT equipment is more susceptible to having its availability affected by highly aggressive techniques, modern OT equipment generally allows for more flexible usage of testing tools.

Two primary factors were observed that could affect the availability of the tested OT devices. Firstly, the throughput of the data sent by the testing tools to the PLCs directly correlated to how much availability was affected. Despite the S7-1200 having similar hardware resources to the ET-200S in terms of capacity, it was more resilient to most of the selected tools. This is most likely due to both a combination of the hardware speed, which is not fully documented in the case of work memory for Siemens PLCs and optimisations provided by more recent firmware. In contrast, despite the Logix5561 having considerably better hardware than both Siemens PLCs, it performed worse during testing with a Ping Flood.

To further demonstrate the effect of network throughput on availability, a second experiment was conducted to determine the capability of these devices to operate appropriately under different network conditions. For this purpose, a custom script was written; it gradually increases the throughput of data being sent to the target devices to determine the thresholds at which each device could perform before observing both a non-negligible increase in latency and packet loss. The results are illustrated in Figure 5.2, and clear distinctions can be observed in how these devices handle different network throughputs. More specifically, for the Siemens ET-200S, an increase in latency can be observed at 400 packets per second (i.e. 25.6 KB/s due to each packet having a size of 64 bytes), and a start of packet loss occurs at 4000 packets per second (i.e. 2.56 MB/s). For the Siemens S7-1200, an increase in latency and a start of packet loss can be observed at 1000 packets per second (i.e. 640 KB/s). For the Allen-Bradley Logix5561, an increase in latency can be observed at 1000 packets per second (i.e. 640 KB/s), and a start of packet loss occurs at 1100 packets per second (i.e. 704 KB/s). For the Windows 7 workstation, no noticeable increase in latency is observed, and a slight increase in packet loss (6%) occurs at 100,000 packets per second (i.e. 6.5 MB/s). The results further reinforce the findings from the experiment detailed in section 5.3.2. These demonstrate that legacy OT (i.e. the Siemens ET-200S) is highly susceptible to disruption during security tests that employ aggressive tools and techniques but that modern OT (i.e. the

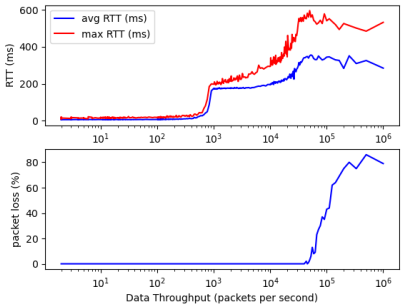

(a) Siemens ET-200S

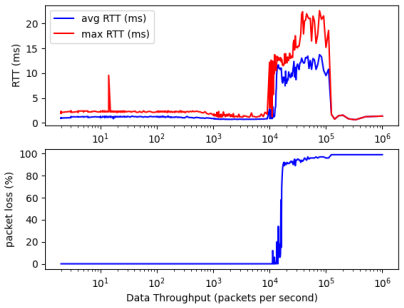

(b) Siemens S7-1200

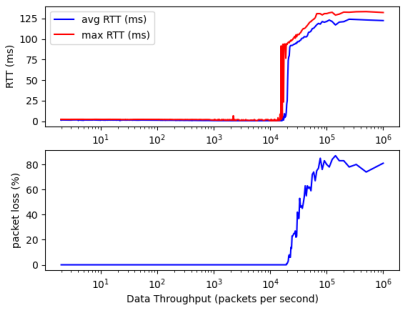

(c) Allen-Bradley Logix5561

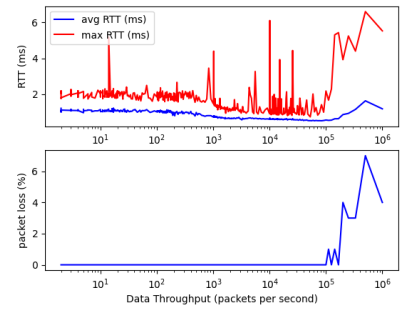

(d) Windows 7 Workstation

Fig. 5.2 Network Stress Test Results

AB logix5561 and Siemens S7-1200), while not as resilient as IT, can be tested with more flexibility.

The second factor that could affect the availability of the tested OT devices was the capability of these devices to process unexpected requests. For example, when vulnerability scripts were used on the ET-200S, it could not process these and resulted in a complete system crash, leading to a total loss of availability and requiring a manual reset. This, therefore, would have a detrimental effect on the environment, causing an adverse impact on the operational process. However, both the S7-1200 and the Logix5561 were able to process these by either handling such requests appropriately or ensuring proper exception handling if an error occurs, allowing for more flexible use of these tools.

Several requirements need to be defined when performing adversary-centric security tests on OT. First, both penetration testing and OT expert knowledge are required to understand precisely how specific tools interact with the system under consideration, as each endpoint will likely react differently to an identical set of tools depending on their hardware and software/firmware. Several factors, therefore, need to be considered concerning this, including determining which protocols the system can process, how errors are handled when encountering unknown requests, etc. The aggressiveness of the tools used during the engagement also

needs to be considered to prevent disruption to the operational process, following availability reduction tolerance. An in-depth understanding is therefore required to understand precisely how specific tools could affect target endpoints which can be provided by both automation and safety engineers. From there, risk quantification can be performed to assess the full scope of the engagement while ensuring that the impact on the operational process is tolerable and that the engagement itself is as comprehensive as possible.

5.4 Conclusion

In this Chapter, the technical differences between Information Technology (IT) and Operational Technology (OT) were analysed specifically within the context of adversary-centric security testing. Considering the technical differences between IT and OT, several challenges were identified when conducting adversary-centric security tests within OT environments. The adversary-centric security testing process can be grouped into phases following the Lockheed Martin and the SANS ICS Cyber Kill Chain. Further analysis showed how the Tactics, Techniques and Procedures used during these phases need to be considered based on what systems or environments are being tested. During the reconnaissance phase, passive techniques, as shown in table 5.1, were found to have little to no impact on the operational process but provided less actionable intelligence for subsequent phases of the CKC. Despite the high probability of causing operational impact if not used properly, especially within OT environments, active reconnaissance techniques, as in table 5.2, were found to return significant actionable information allowing for more depth of testing to be made. The weaponisation stage of the CKC was identified as being closely correlated to the impact goals of adversaries, which can differ significantly between IT and OT targets, and are summarised in table 5.3. Further phases of the CKC found that the TTPs used during these were often similar in execution, albeit modified to suit targeted endpoints.

While commonly used tools for IT-centric engagements may not have any noticeable effect in these environments, it is possible that they can adversely disrupt the operational process within OT environments. This was validated by deploying tools with varying degrees of aggressiveness on industrial control systems. Findings from this exercise identified two factors that could adversely affect the operational process through a reduction or loss in availability. First, the network throughput of active tools was directly correlated to a loss of availability, the rate of which is unique to the system under consideration based on hardware and software capabilities. Second, the use of unexpected techniques such as vulnerability scans and scripts resulted in operational impact depending on the targets' capabilities for processing and handling errors.

Despite current approaches that limit the use of adversary-centric security testing tools to strictly passive ones during assessment engagements, employing active tools is possible subject to the resilience of the systems against more aggressive techniques, as demonstrated in section 5.3. While existing frameworks for adversary emulation for security testing exist, such as MITRE's Adversary Emulation Plans which provide techniques and tools for emulating specific threat actors [177], these do not take into consideration the safety and operational risks that security testing can present to OT environments. Chapter 6 aims to identify how to comprehensively quantify the risk that active adversary-centric security testing techniques have on ICS/OT and to allow for better scoping of these engagements. This will minimise the risks that these techniques present to safety and the operational process while ensuring the full depth of such an engagement as part of the overall cyber risk assessment life cycle.

Chapter 6

Risk-Based Safety Scoping of Adversary-Centric Security Testing on Operational Technology

Chapter 5 presented the current challenges of conducting adversary-centric security tests within ICS/OT environments. Because of their critical nature and the design philosophies regarding ICS/OT, additional risks need to be considered in order to prevent impact to safety and the operational process during these engagements. In recent years, however, newer product lines from OT vendors, such as Siemens [247] or Allen-Bradley [4], have seen an increase in performance, allowing for more flexibility during adversary-centric security testing. Identifying and understanding the risk that tools and techniques used during such engagements still needs to be undertaken so that scoping of such tests can consider these risks so as not to disrupt the operational process. This chapter provides a methodology for the identification and quantification of safety and operational risk during security testing and proposes a framework to scope adversary-centric security tests as a means of maximising the depth-of-testing while minimising safety and operational risk.

The existing safety risks confirmed by Chapter 5 which identified that the safety-critical nature of ICS/OT environments requires unique scoping of adversary-centric security tests so that safety risks can be minimised while ensuring that depth-of-testing is maximised. While testing multiple OT devices from vendors, including Siemens and Allen-Bradley, two main factors of using adversary-centric security testing tools were identified that could cause a reduction in availability or integrity and disrupt the operation process. Firstly, high network traffic generated by these tools could cause an increase in latency or an observable loss in transmitted packets. Secondly, the data being sent to the target could cause additional overhead on its resources, resulting in either a reduction in availability through resource

exhaustion or total loss of availability due to some data not being processed appropriately and causing a system crash. While some of the tools used during testing consistently resulted in a severe loss of availability, this was not the case for a majority of them; demonstrating that adversary-centric security testing within ICS/OT environments is indeed possible if the effects of the tools and techniques used are understood and taken into consideration during scoping of engagements. This Chapter provides a methodology for identifying and quantifying the safety and operational risks of conducting adversary-centric security tests within ICS/OT environments. These methodologies are then used for the creation of a framework to aid in the scoping of these.

6.1 Identifying Safety and Operational Risks of Adversary-Centric Security Testing on ICS/OT

6.1.1 Identifying hazards with (C)HAZOP

Derived from the well-established Hazard and Operability (HAZOP) study [146], a Control Hazard and Operability (CHAZOP) study provides a comprehensive framework for reviewing controllability, safety and operability issues during the implementation of ICS/OT [160]. The objective of such a study is to understand and assess hazards that could cause a loss of safety or a disruption to the operational process, which is the first step in quantifying the safety and operational risks of conducting adversary-centric security tests within ICS/OT environments. While several methodologies exist for identifying hazards, HAZOP was found suitable for the identification of hazards caused by adversary-centric security tests within ICS/OT environments due to its applicability for identifying both safety and operational hazards and its widespread application across several domains, including manufacturing, engineering and CNI [65]. Additionally, while the Institution of Chemical Engineers acknowledge that certain factors such as no prior design review; inappropriate, incompetent or too many team members; lack of operational experience; defensive designers; and arrogant project managers can reduce the effectiveness of (C)HAZOP studies, if executed correctly, these type of studies allow for effective and cost-efficient qualitative risk assessment [283].

When applying these studies to adversary-centric security testing, the terminology used is similar to that used in HAZOP studies with additional context. These are as follows:

- Node: The specific location in the process for which deviations can occur (for example: heater, liquid tank, mixers).

Guideword	Definition	Example
NO or NOT	Complete Negation of the Intention	No Flow; No Communication; No Pressure
MORE and LESS	Quantitative Increase or Decrease	More/Less Flow; Less Communication; More/Less Pressure
AS WELL	Qualitative Increase	Intended Valve Close As Well As Unintended Valves
PART OF	Qualitative Decrease	Part Of Intended Valves Closing
REVERSE	Opposite of the intention	Reverse Flow; Reverse Direction
OTHER THAN	Complete Substitution	Other Than X Chemical

Table 6.1 (C)HAZOP guidewords

- Parameter: The parameter for the condition(s) of the process (for example: temperature, level, flow, pressure).
- Intent: How the node is designed to operate under normal conditions.
- Guidewords: Terms when considered with one or more parameters that form a hypothetical deviation for risk consideration (i.e GUIDEWORD + PARAMETER = DEVIATION).
- Deviations: Events that lead to a partial or total disruption of the operational process.
- Causes: The combination of the events that cause deviation.
- Consequences: The outcome derived from the causes that could lead to operational impact or loss of safety.
- Actions: Actions that can be taken to mitigate the identified risk(s).

The methodology for applying a (C)HAZOP study in the context of adversary-centric security testing is depicted in Figure 6.1. As opposed to HAZOP, (C)HAZOP focuses on hardware and software design of ICS/OT rather than vessels and pipes. Any system related to safety or operation functions should be considered during the study. For each of the identified endpoints, the following must be considered to comprehensively understand the risk that these face:

- The functionality of the system.
- All the dependencies of the system.
- Segregation and redundancy deployments.
- Application of the guidewords from Table 6.1.

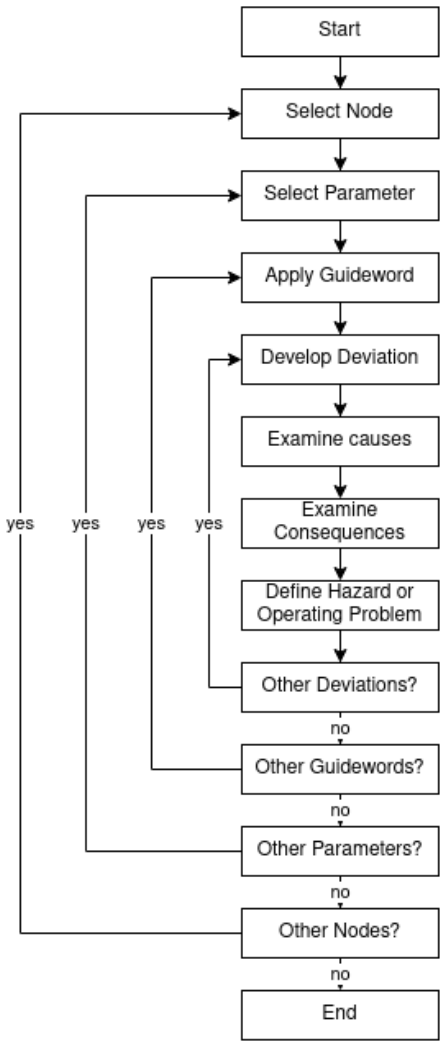

Fig. 6.1 HAZOP Methodology

While the methodology in Figure 6.1 is generally only applied to endpoints used for safety-related functions, for adversary-centric security testing, parameters that affect the operational process should also be considered. When applying (C)HAZOP to the risk assessment process, it is essential to ensure the full coverage of documents is considered and should include the following:

- User Requirement Specification and Detailed Functional Specification documents.
- Piping and Instrumentation (P&I) Diagrams.
- Network Diagrams.
- System hardware configuration documents.
- Power and wiring documents.
- Channel/loop diagrams.
- System malfunction fail-safes.

To demonstrate an example application of (C)HAZOP for identifying hazards during an adversary-centric security test, this process has been applied to a scenario engineered within the Lancaster University ICS testbed; previously used for the creation of the attack scenarios described in Chapter 3. The scenario consists of an operational process to manually control the water levels of a tank through an HMI panel and contains the following elements:

- Siemens SIMATIC ET-200S (physical device): sends data to the HMI, receives commands from the HMI, receives data from the water tank sensor, and sends commands to the tank pump and release valve.
- Siemens TP1500 Basic PN HMI (physical device): displays water tank levels, receives data from PLC, sends open/close commands for both the tank pump and the release valve of the water tank to the PLC.
- Water Tank (virtualised): Container for water storage.
- Water Tank Pump (virtualised): turns on and off to increase water level in the water tank.
- Water Tank Release Valve (virtualised): opens and closes to decrease water level in the water tank.

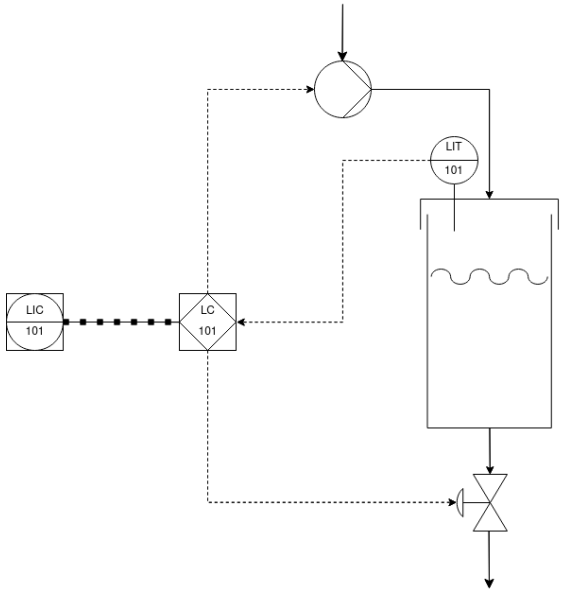

Fig. 6.2 P&I Diagram of Water Tank Scenario

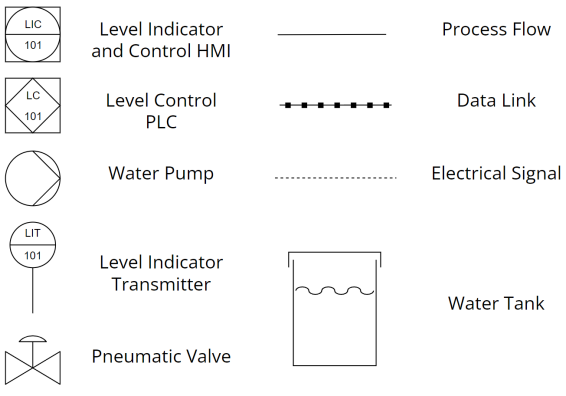

Fig. 6.3 P&I Diagram Symbols

Parameter	Guideword	Deviation	Causes	Consequences
Tank Water Level	More	More Water Level	Pump On and Water Level High	Tank Overflow
Pump	No	No Pump (De)activation	Pump unresponsive	Disruption to Operational Process
Release Valve	No	No Valve (De)activation	Valve unresponsive	Disruption to Operational Process
PLC	No	No PLC Communication	PLC Resource Overload; PLC Crash	No Control of Pump and Release Valve
PLC	Less/Late	Less/Late PLC Communication	PLC Resource Overload; Network Congestion	Limited Control of Pump and Release Valve
HMI	No	No HMI Communication	HMI Resource Overload; HMI Crash	No Control of PLC, Pump and Release Valve
HMI	Less/Late	Less/Late HMI Communication	HMI Resource Overload: Network Congestion	Limited Control of PLC, Pump and Release Valve

Table 6.2 (C)HAZOP Output for Water Tank Scenario

- Water Tank Sensor (virtualised): sends water level data to PLC.

Figure 6.2, of which its symbols are detailed in Figure 6.3, represents an ANSI/ISA-5.1-2009 [5] and ISO 14617-6:2002 [125] compliant P&I Diagram of the scenario developed within the ICS testbed. Despite the scenario being simple in concept, it accurately depicts, at a reduced scale, the potential hazards possible within real-world industrial processes. By applying a (C)HAZOP methodology, several safety and process hazards can be identified and are provided in Table 6.2.

6.1.2 Establishing Risk Events and Causes with FTA

Following the identification of hazards using (C)HAZOP, a Fault Tree Analysis (FTA) can be conducted to further decompose hazards into their causes. While (C)HAZOP can also be used to qualitatively identify the causes of hazards, FTA is used to provide further depth to this by identifying the relationship between different events that could lead to the cause of a major hazard. This analysis adopts a top-down approach where hazards are broken down into possible causes. Each of these causes is then decomposed until a set of "basic events" is established, for which their risk can be calculated. The components of a Fault-Tree Diagram (FTD) are defined within IEC 61025 [113] and are as follows:

- Gates: Symbols (see Figure 6.4) showing the logical relationship between a cause and a consequence. Static gates do not depend on the order of occurrence whereas dynamic gates do.

- Events: Symbols (see Figure 6.4) describing failure states, system states, or events within an even chain.

In order to fully develop a fault tree, a thorough understanding of the cause and effect relationships between a hazard and its subsequent causes is required and can be provided by both safety and ICS engineers. Following a pragmatic methodology, causes need to be determined based on their possibility of occurring during adversary-centric security testing.

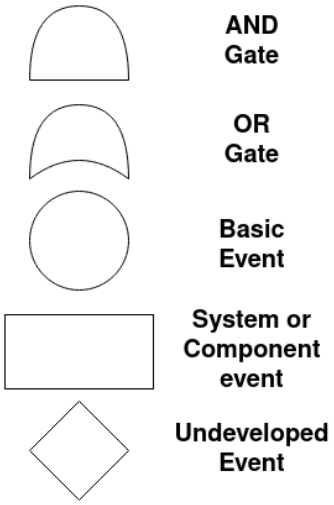

Fig. 6.4 Example Symbols used for Fault Tree Analysis

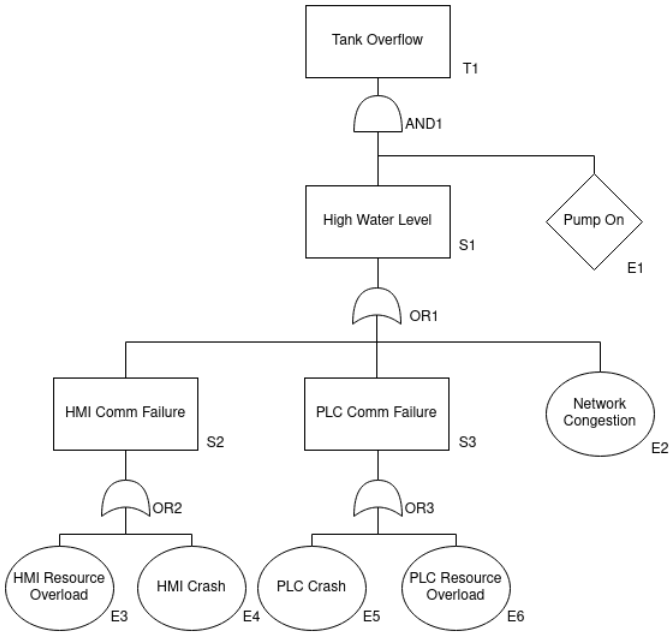

Fig. 6.5 Fault Tree Diagram for Tank Overflow Hazard

As opposed to traditional safety risk assessment, this excludes failure mode risks such as power failure.

Continuing with the example provided by the scenario described in Figure 6.2, Figure 6.5 was developed following an FTA for the Tank Overflow hazard that was identified during the preceding (C)HAZOP study. While there exist safety hazard events such as power supply failures or mechanical valve failures that could lead to this hazard scenario and should be developed within a traditional HAZOP study because this study focuses on the effects of adversary-centric security testing on safety and the operational process, these have been excluded from the final FTD.

Once an FTD has been generated, the minimal cut sets (MCSs) for this can be deduced. These sets are the unique combination of basic events from the FTD that can lead the top event to occur, such as the water tank overflowing from Figure 6.5. For this, top events are denoted as T, system events are denoted as S, and basic events are denoted as E. When determining the MCSs for an identified hazard, OR gates produce additional cut sets, whereas AND gates make the cut sets more complex. For example, to begin developing the MCSs for the tank overflow scenario, AND1 immediately below T1 can be listed as the following expression:

$$T1 = S1 \wedge E1$$

The expansion of $S1 = S2 \vee S3 \vee E2$ leads to the following:

$$S2 \wedge E1$$
$$S3 \wedge E1$$
$$E2 \wedge E1$$

Substituting for $S2 = E3 \vee E4$; and $S3 = E5 \vee E6$ results in the following MCSs (denoted C_i):

$$\begin{aligned} C_1 &= \{E6, E1\} \\ C_2 &= \{E5, E1\} \\ C_3 &= \{E4, E1\} \\ C_4 &= \{E3, E1\} \\ C_5 &= \{E2, E1\} \end{aligned} \tag{6.1}$$

While the MCS provided in the list of sets 6.1 does not require further reduction, more complex cut sets can be reduced by removing redundant events or sets through the idempotence or absorption rule, for example. Because of the complexity of some systems, this can

result in MCSs containing several thousand cut sets. Therefore, truncation can be used to remove cut sets that are believed to contribute negligibly to the top event occurring, which can be determined through traditional safety risk assessment. Additionally, if available, FTA software can also be used to automate the creation of Fault Trees and calculation of MCSs.

6.2 Quantifying Safety and Operational Risks of Adversary-Centric Security Testing on ICS/OT

Once safety and operational hazards have been identified, these can be evaluated to determine the risk of conducting adversary-centric security tests within ICS/OT environments. As such, by understanding and assessing these risks, strategies can be formulated to appropriately scope these engagements and ensure their completeness while mitigating the potential for operational disruption and loss of safety. Safety and operational risk is commonly defined as a product of likelihood and impact, where likelihood refers to the probability of a risk event occurring and impact refers to the severity of the consequences when a risk event occurs. Due to the operational nature of ICS/OT environments, the impact of events can be represented through either monetary cost (for hazards leading to disruption of the operational process) or injuries/deaths (for hazards leading to a loss of safety). Both expert estimation and historical data can be used to calculate the impact of an event occurring in their respective environments. The following subsections describe the methodology for quantifying the likelihood of hazards occurring by calculating the probability of the respective basic events occurring based on safety and operational failures. This can subsequently be used in the overall risk quantification of identified hazards. All data and scripts used for quantification of risk have been made publicly available on GitHub [262].

6.2.1 Cut Set Probability

As part of the evaluation of a Fault Tree (discussed in section 6.1.2), the probability of top events can be calculated based on the probability of the bottom events occurring. Because the fault tree of real systems commonly contains recurring basic events, this evaluation can be done using derived MCSs. For example, given the MCSs determined for the scenario described in Figure 6.2, the top event (Tank Overflow) can be expressed as the following boolean expression:

$$\begin{aligned}\text{TankOverflow} =&(E2 \wedge E1) \vee (E3 \wedge E1) \vee (E4 \wedge E1) \\ &\vee (E5 \wedge E1) \vee (E6 \wedge E1)\end{aligned}$$

As such, the probability for the top event occurring can be expressed as follows:

$$\begin{aligned} P(\text{TankOverflow}) =& P((E2 \wedge E1) \vee (E3 \wedge E1) \vee \\ & (E4 \wedge E1) \vee (E5 \wedge E1) \vee (E6 \wedge E1)) \end{aligned}$$

As each MCS is capable of causing the top event, their likelihood to cause the top event is therefore cumulative. However, each MCS may not be mutually exclusive (i.e. non-disjoint) since these can contain the same basic event. Due to the rule of addition, the probability of each MCS occurring will be greater than or equal to the probability of the top event occurring. For example, E1, E2 and E3 could coincide, satisfying the first two MCSs. Because of this, the upper-bound of the probability of the tank overflowing scenario can be defined as:

$$\begin{aligned} P(\text{TankOverflow}) \leq & P(E2 \wedge E1) + P(E3 \wedge E1) + \\ & P(E4 \wedge E1) + P(E5 \wedge E1) + \\ & P(E6 \wedge E1) \end{aligned} \tag{6.2}$$

While using term combination does increase the accuracy of the probability of a top event occurring, the resulting formula for this becomes exponentially more complex the more MCSs are present, which is especially common for large fault trees. Furthermore, the subsequent combination of terms within a derived formula, otherwise known as the "rare event contribution", contribute significantly less to the probability of the top event occurring than the first terms established from the FTA. Therefore the approximation provided in equation 6.2 can be deemed adequately accurate for subsequent risk analysis as it provides an upper bound for the probability of an event occurring.

Because the events contained within an MCS are independent, as per the definition of a basic event, the final upper-bound probability of the tank overflowing can be further decomposed as follows:

$$\begin{aligned} P(\text{TankOverflow}) \leq & P(E2) \times P(E1) + P(E3) \times P(E1) + P(E4) \\ & \times P(E1) + P(E5) \times P(E1) + P(E6) \times P(E1) \end{aligned} \tag{6.3}$$

The following formula can, therefore, be used to calculate the upper-bound of the probability of a safety or operational hazard occurring during an adversary-centric security test using MCSs:

$$P(\text{TopEvent}) \leq \sum_{j=1}^{k} \left[\prod_{E \in C_j} P(E) \right] \tag{6.4}$$

Where E is a basic event belonging to a minimal cut set C_j and k is the total amount of MCSs.

6.2.2 Basic Event Probability

To provide further granularity in determining the risk of top events, the probability of the basic events belonging to the MCSs of an associated top event needs to be calculated. Previous work identified two contributors to basic events leading to safety and operational hazards during adversary-centric security tests [265]. The first is due to excessive data throughput of tools being used (named Network-Caused Basic Events), and the second is due to the contents of the data. The second event type can be further decomposed into two sub-categories: data that causes excessive overhead (named Resource Exhaustion Basic Events) and data that, when processed by an industrial device, results in a system crash or error (named Incompatible Data Basic Events). From the FTD illustrated in Figure 6.5 for the water tank scenario, basic event E2 can be categorised as a Network-caused basic event, basic events E3 and E6 can be categorised as resource-exhaustion basic events, and basic events E4 and E5 can be categorised as Incompatible data basic events.

Network-Caused Basic Events

To obtain accurate data on how the data throughput of tools or techniques used during an adversary-centric security test on ICS/OT could have an adverse effect on the operational process of an industrial environment, a network stress test can be performed on the target endpoints; done within a testing environment such as a testbed to prevent impact to the operational process. By gradually increasing the amount of data being sent to the target, the endpoint's capability of responding to high network traffic can be assessed and thus, the throughput of data at which an increase in latency or a loss of packets would lead to disruption to the operational process, can be determined.

Continuing with the example provided in the P&I Diagram from Figure 6.2, both the Siemens HMI and PLC need to be tested to determine the limits of their packet buffer and the effect of high throughput tools and techniques on these. For this, a custom script was created to simulate network traffic using ICMP ping packets with decreasing delay between packets to determine the behaviour of these devices with different network throughputs. The results from this test can be found in Figure 6.6 and Figure 6.7. For the ET-200S, a considerable increase in latency can be observed at around 400 packets per second, equating to approximately 25.6 KB/s (due to each packet used during the test being 64 bytes in size). However, no packet loss is observed until around 40000 packets per second, which equates

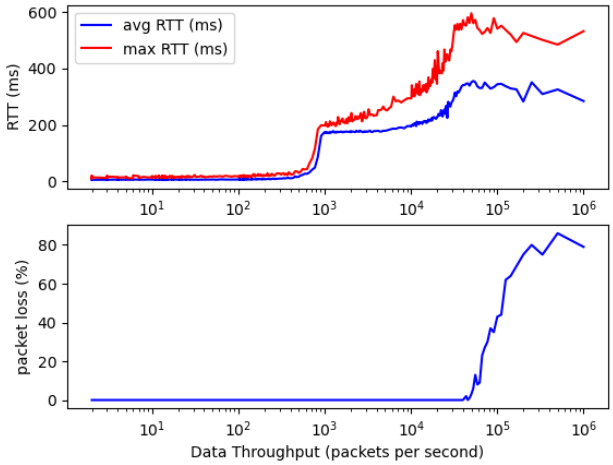

Fig. 6.6 ET200S Data Throughput Test Results

to a throughput of approximately 2.56 MB/s. However, results from testing the HMI show a near-total packet loss at 1000 packets per second with no increase in latency prior to this. From this data, in order to prevent any disruption to the operational process, all tools and techniques used during an adversary-centric security test within this environment would need to output data at a rate of less than 25.6 KB/s.

While it is possible to determine a maximum tolerable throughput for adversary-centric security testing activities, inherent network jitter can contribute to additional risk in environments with strict timeliness requirements, such as CNI, and therefore must be determined. While some tools might seem safe for use within certain environments due to their low inherent network throughput, they may cause additional jitter leading to the possibility of reduced availability and therefore must also be considered. Several works have attempted to estimate the distribution of network jitter with varying results. For example, Karakas determined that network jitter distribution can mostly be fitted to a lognormal distribution if no additional factors, such as firewall ruling, contribute to network delay [138]. However, Mozhaiev et al. claim that random jitter is best fitted to a Gaussian Distribution [183] and Daniel et al. describe network jitter as fitting a Laplacian distribution [51]. This disparity in distribution fits is mainly attributed to the causes of network jitter, such as random noise, crosstalk from signals, the effect of dispersion from signal propagation, or resistance mismatch, which all affect the distribution of network jitter differently. As such, the distribution of network jitter is dependent on the environment itself and therefore needs to be determined for each environment which is planned to be tested.

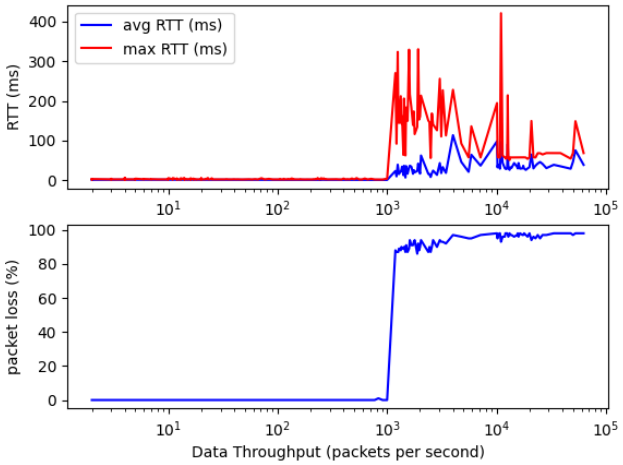

Fig. 6.7 Siemens HMI Data Throughput Test Results

To determine the network jitter distribution within the water tank scenario, data was collected on the latency of both the ET-200S and the HMI while these were continuously receiving 25.6 KB/s of data for 15 minutes. The results of this experiment can be found in figures 6.8 and 6.9. Using the python library `distfit`, the Residual Sum of Squares (RSS) was calculated for the best fitting distributions, which were the lognormal distribution (RSS=0.048644 for the ET-200S data and RSS=2.496344 for the HMI data) and the Generalised Extreme Value (GEV) distribution (RSS=0.04456 for the ET-200S data and RSS=2.646364 for the HMI data). All results from using `distfit` to calculate the various fitness scores of distributions can be found in Tables 6.3 and 6.4. While the GEV distribution was deemed to be a better fit for the jitter distribution of the ET-200S, the shape parameter of the GEV distribution for the HMI was negative, suggesting that this distribution has an upper limit as per its definition when using negative shape parameters. As network jitter can cause, in extreme cases, high latency values leading to packet loss, the GEV distribution was therefore rejected, and the lognormal distribution was selected as the most appropriate distribution fit for the water tank scenario's network jitter.

As such, the following 3-parameter formula can be used to calculate the probability density function for an endpoint's latency based on jitter within the water tank scenario described in Figure 6.2; the curve for these is illustrated in figures 6.8 and 6.9:

$$f(x;m;s;\theta) = \frac{1}{(x-\theta)s\sqrt{2\pi}} exp\left(-\frac{(\ln(\frac{x-\theta}{m})^2}{2s^2}\right) \qquad (6.5)$$
$$x > \theta; m, s > 0$$

ranking	distr	score	LLE	loc	scale	arg
0	genextreme	0.04456	NaN	4.124234	0.715471	(0.019880619231587585,)
1	lognorm	0.048644	NaN	1.341275	3.059706	(0.2670700729681621,)
2	gamma	0.055465	NaN	2.094081	0.30557	(7.918591272636394,)
3	beta	0.055789	NaN	2.127674	1787830.197773	(7.6293048670266215, 5714159.193619767)
4	t	0.111718	NaN	4.43861	0.710655	(6.171637809106455,)
5	dweibull	0.117561	NaN	4.46801	0.712976	(1.2178260986783505,)
6	norm	0.135643	NaN	4.513763	0.899412	()
7	loggamma	0.146162	NaN	-354.035654	45.900016	(2469.654144604755,)
8	expon	0.982031	NaN	2.653	1.8607632	()
9	uniform	1.326598	NaN	2.653	10.4656	()
10	pareto	1.384878	NaN	0.001607	2.651393	(1.9463309032892653,)

Table 6.3 Best Fit Results using `distfit` for PLC Network Distribution at 400 packets per second

ranking	distr	score	LLE	loc	scale	arg
0	lognorm	2.496344	NaN	0.686484	0.208651	(0.519992978790028,)
1	genextreme	2.646364	NaN	0.861475	0.084229	(-0.16375062691567044,)
2	beta	2.824626	NaN	0.714093	1592192285724.914062	(2.8199955923958058, 21167823283185.176)
3	dweibull	5.24911	NaN	0.886558	0.098125	(1.0069936788017444,)
4	t	7.237047	NaN	0.907803	0.101204	(5.579147878141951,)
5	norm	9.109383	NaN	0.926964	0.168039	()
6	loggamma	9.893515	NaN	-51.624706	7.135857	(1579.2097368296409,)
7	expon	16.320509	NaN	0.719	0.207964	()
8	pareto	17.57609	NaN	-1085520.353355	1085521.072355	(5560848.1067601815,)
9	gamma	33.248012	NaN	0.719	0.47821	(0.1306685429804617,)
10	uniform	39.455348	NaN	0.719	3.175	()

Table 6.4 Best Fit Results using `distfit` for HMI Network Distribution at 400 packets per second

where:

- x is a given RTT in milliseconds.
- θ is the location parameter of the distribution.
- m is the scale parameter of the distribution.
- s is the shape parameter of the distribution.

It is worth nothing that the probability density function of both figures 6.8 and 6.9 are represented through a histogram, meaning that the probability for a given RTT range is defined as the following:

$$P(bar_{min} < X < bar_{max}) = (bar_{max} - bar_{min}) \times bar_{height}$$

Using the parameters derived from either interpolation techniques such as curve fitting or their respective formulas, the cumulative distribution function (CDF) of an endpoint's latency

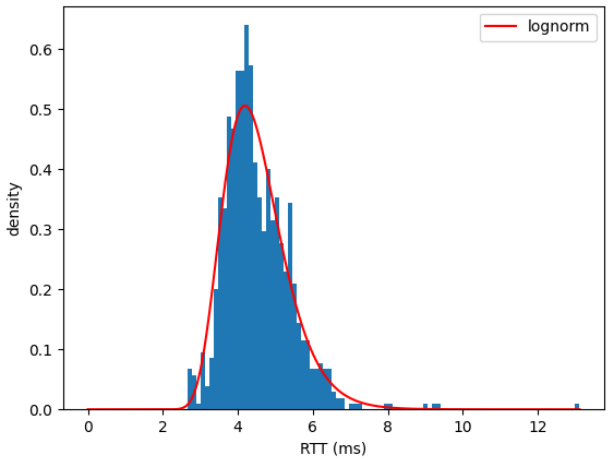

Fig. 6.8 Lognormal Distribution Curve Fit of ET-200S Network Jitter

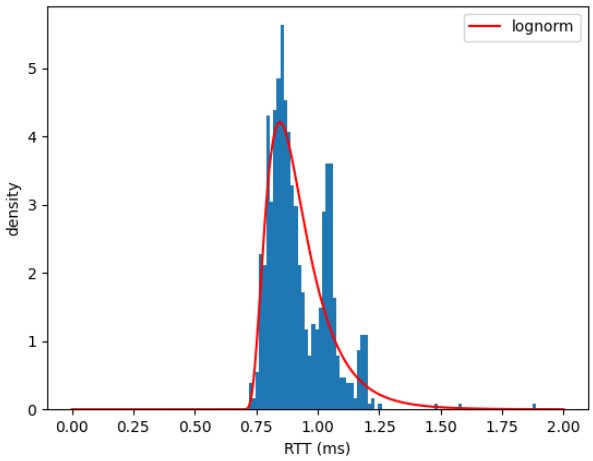

Fig. 6.9 Lognormal Distribution Curve Fit of Siemens HMI Network Jitter

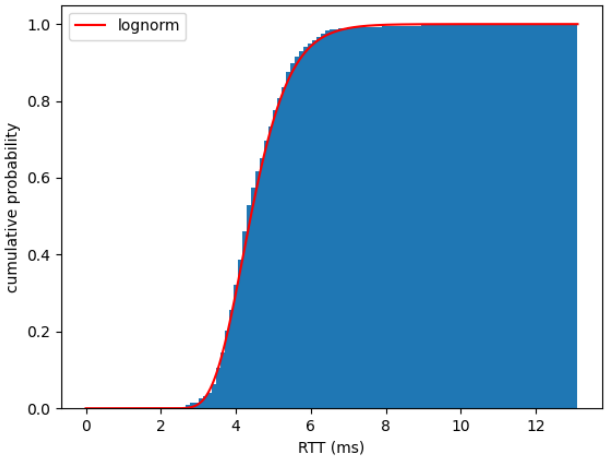

Fig. 6.10 Jitter Cumulative Probability for ET-200S during 400 Packets (of 64 Bytes) per Second Test

can be used to determine the probability of an adversary-centric security tool or technique's throughput causing undesirable latency and affecting the operational process.

$$F_x(x;m,s,\theta) = \Phi\left(\frac{ln(\frac{x-\theta}{m})}{s}\right) \qquad x \geq \theta; m,s > 0 \tag{6.6}$$

where:

- $\Phi(x)$ is the cumulative distribution function of the standard normal distribution ($\Phi(x) = \int_{-\infty}^{x} \frac{exp(\frac{-x^2}{2})}{\sqrt{2\pi}}$)

Using the CDF derived from data for the Siemens HMI and the ET-200S, the probability of the latency exceeding a tolerable value can be estimated, which must be determined, specific to the environment and endpoints being tested, by safety and ICS engineers. Example tolerable latency values for the HMI and PLC from the scenario described in Figure 6.2 were arbitrarily determined to be 3ms and 15ms, respectively. The probabilities of the latency of these endpoints exceeding these values while receiving 25.6KB/s of data from adversary-centric security testing tools and techniques were determined as follows using the CDF from equation 6.6:

$$P(\text{HMI\_latency} > 3ms) = 3.25 \times 10^{-7}$$

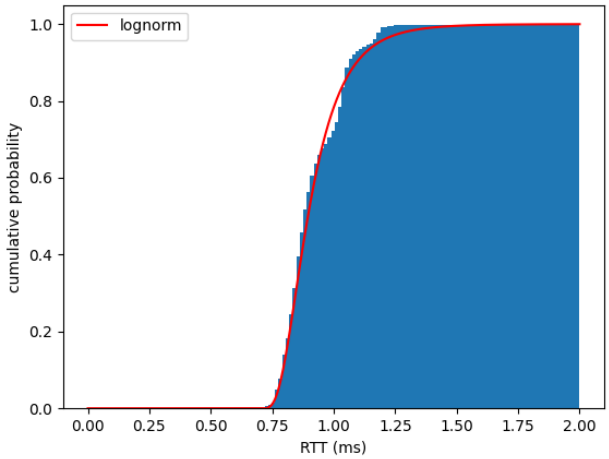

Fig. 6.11 Jitter Cumulative Probability for Siemens HMI during 400 Packets (of 64 Bytes) per Second Test

$$P(\text{PLC\_latency} > 15ms) = 1.06 \times 10^{-8}$$

These probabilities can subsequently be used in determining the probability of Network-Cause Basic Events that contribute to a top event occurring in the scenario described in Figure 6.2 as discussed in Section 6.2.1.

Resource Exhaustion Basic Events

While data throughput exceeding tolerable ranges can disrupt the operational process, some tools and techniques employed during an adversary-centric security test may cause similar disruption, due to endpoint resource exhaustion, without exceeding these ranges. Because of this, testing also needs to be done to determine if any tools or techniques planned to be employed throughout an adversary-centric security test could cause disruption due to resource exhaustion. If data from previous engagements is unavailable, this needs to be obtained through experimentation in a testing environment such as a testbed. Expert opinion can aid in estimating the effect of tools and techniques; however, testing is required for tools or techniques that have an unknown effect on an endpoint's resources.

For example, port scanning is a commonly-used technique employed during adversary-centric security tests to discover open ports on an endpoint. By identifying these, the devices' services can be deduced and tested further for vulnerabilities. Nmap is a popular tool used for port-scanning, allowing for different scan options. As such, a comprehensive test of all these options must be done to determine which of these presents the least risk, if any, for

port-scanning ICS/OT. The following port scan options were therefore tested on the ET-200S within the context of the water tank scenario described in Figure 6.2:

- Idle (control test),
- TCP SYN scan (uses SYN packet but does not complete full TCP handshake),
- TCP Connect scan (full TCP handshake),
- UDP scan (only used to determine open UDP ports),
- SCTP INIT scan (uses the SCTP protocol over TCP/UDP),
- TCP NULL scan (no flags set),
- TCP FIN scan (TCP FIN flag set only),
- TCP Xmas scan (TCP FIN, PSH, and URG flags set),
- TCP ACK scan (ACK flag set),
- TCP Window scan (examines TCP Window field of the returned RST packets),
- TCP Maimon scan (TCP FIN and ACK flags set).

By acquiring data on PLC CPU execution time with no additional load, a baseline can be determined to identify abnormally high increases in execution time, which could disrupt normal functions. The results from running these scanning options continuously for 15 minutes on the ET-200S have been summarised into boxplots, illustrated in Figure 6.12. These boxplots allow us to identify the non-outlier minimum, non-outlier maximum, median, first quartile, third quartile and outliers of CPU execution times for each scan option.

To obtain additional precision on how these scan options can impact the operational process, data on the effect of these on an endpoint's network response time can be used. As such, Figure 6.13 summarises the results of testing the ET-200S' latency when being scanned continuously for 15 minutes with the same scan options as Figure 6.12.

From this data, four scan options can be identified with certitude as being high risk due to both the observed CPU execution times and latency of these tests being considerably higher than that of the control test. These include the Window, TCP ACK, TCP FIN, and TCP SYN scans. This is because these scan types cause additional load on endpoints to increase the scan's stealth or speed. For SYN scans, as an example, the speed of the scan is increased because the TCP three-way handshake remains incomplete. However, due to this,

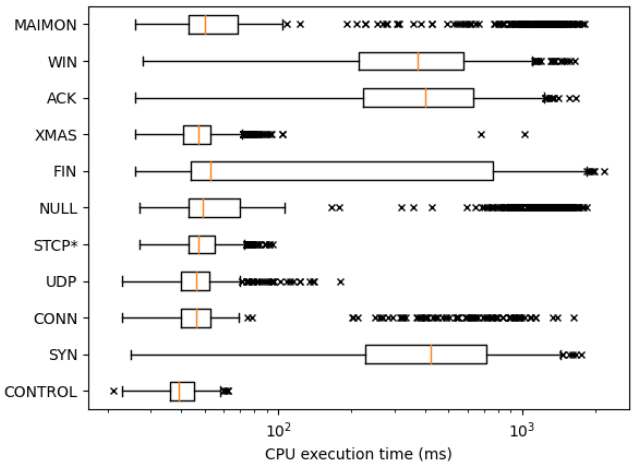

Fig. 6.12 ET-200S CPU Execution Times with Nmap Scan Options

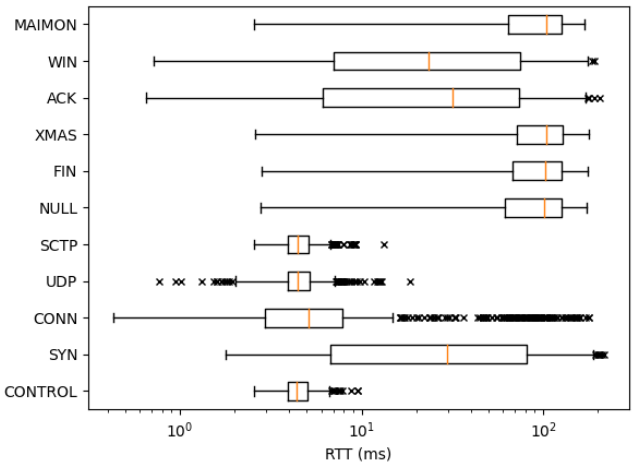

Fig. 6.13 ET-200S Latency with Nmap Scan Options

the endpoint continuously allocates resources for incoming TCP connections, which never occur, leading to the potential of a SYN flood and, consequently, the potential to disrupt the operational process.

Despite the Xmas scan, the Maimon scan and the NULL scan not resulting in high CPU execution times, a non-negligible increase in latency was observed. This is most likely the result of these tests being conducted with default scan speeds (no initial scan delay and dynamic parallelism). Therefore, as a means of reducing the risk of latency issues causing disruption to the operational process, less aggressive scan speeds can be used, such as the polite (initial scan delay of 400ms and max parallelism of 1) or sneaky (initial scan delay of 15 000ms and max parallelism of 1) options. However, other scan options, discussed subsequently, with normal scan speeds, present considerably less risk and should be favoured over the Xmas, Maimon and NULL scan options.

While the SCTP scan resulted in a negligible increase in CPU execution time and latency for the ET-200S, the expected scan results were not returned. Despite port 102 (S7COMM) being open, the SCTP scan identified it as closed. This is due to the PLC not supporting this specific protocol and therefore not replying with appropriate data for identifying open ports. This scan option is therefore not recommended for use on this PLC specifically.

The remaining scan options, which include the UDP scan and TCP Connect scan, resulted in both a negligible deviation of CPU execution time and acceptable increases in latency while also returning correct information on open ports. The UDP scan is unique because it is the only scan option available for identifying open UDP ports. Fortunately, using this option on the ET-200S does not result in any significant increase in CPU execution time and can therefore be considered safe to use depending on established risk tolerance. While causing some increase in CPU execution time, the TCP Connect scan causes less disruption than the other tested scan options. However, these increases in CPU execution time are expected as any additional load on the PLC will lead to increased CPU execution time regardless of the task. Furthermore, these outliers ($execution\ time > 100ms$ for the TCP Connect scan) only consist of 9% of total registered execution times. Additionally, the speed of the scan can be configured to reduce increases in latency and further reduce risk to the operational process. Therefore, if within established tolerable ranges, both the UDP and TCP Connect scans are the safest scan options for use on the ET-200S.

Due to the Siemens HMI not having diagnostic capabilities, acquiring data on resource usage is more challenging than for the PLC. Despite this, measuring network latency alone, while not as accurate as measuring both this and CPU execution time, provides sufficient estimation of the effects of different port scanning options due to changes in latency, in most cases, correlating with resource usage as observed when testing the ET-200S. The results of

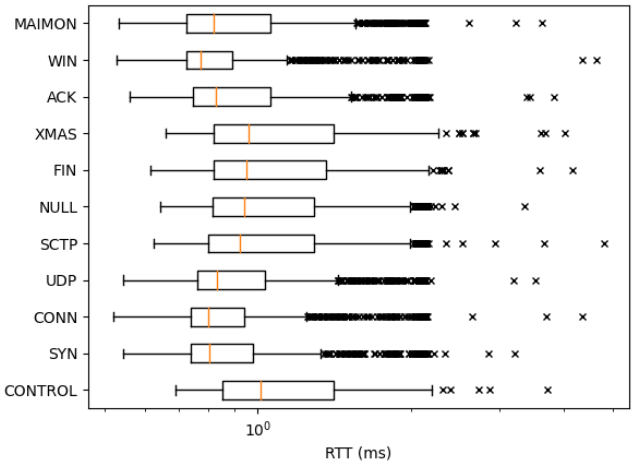

Fig. 6.14 Siemens HMI Latency with Nmap Scan Options

running these scan options continuously for 15 minutes on the Siemens HMI are summarised in Figure 6.14.

Because the HMI runs WinCC on top of a Windows Operating System (Windows Embedded Compact) and uses better hardware than the PLC, it is considerably more resilient to the different scan options available with Nmap. As seen in Figure 6.14, all of the scan options used on the HMI returned similar results and no considerable increase in latency was observed as opposed to the results from conducting the same test on the ET-200S. Therefore, most scanning options can be considered safe for use on the HMI. However, if additional risk reduction is required, this can be done by using less aggressive scan speeds, similar to the ET-200S.

Incompatible Data Basic Events

Despite basic events caused by incompatible data presenting the most danger to the operational process, identifying and quantifying these is relatively simple. During testing for resource exhaustion basic events, any tool or technique which consistently results in the failure of data integrity or the failure of exception handling needs to be identified and marked during scoping to prevent the use of these during an engagement.

During testing of the ET-200S used for the water tank scenario described in Figure 6.2, three open-source and commercial tools were identified as affecting the PLC's behaviour to the point of disrupting the operational process: Nmap (service and version enumeration), Nessus, and OpenVAS. While running these, the PLC would enter an error state, disrupting

all communication to the HMI and actuators and requiring both a complete power cycle and a master reset to restore the PLC to a working state.

Despite Nessus and OpenVAS initially causing total disruption to the operational process, a change in the configuration of the ET-200S was identified to prevent the PLC from entering an error state. By loading a programming error Organisation Block (designated OB121 in the TIA Portal, used for programming Siemens PLCs) into the CPU load memory, subsequent scans using Nessus and OpenVAS did not result in any error state occurring. However, further testing following the methodologies described in Sections 6.2.2 and 6.2.2 would still need to be undertaken to determine network and resource-related risks when using these tools.

While loading OB121 into the CPU load memory resolved error states caused by using Nessus and OpenVAS, subsequent scans using Nmap's service and version enumeration module still resulted in the PLC entering an error state. Upon analysis of the packets sent by Nmap prior to the PLC crashing, Nmap attempts an RDP Negotiation Request with the PLC as part of an RDP Connection Request Protocol Data Unit. Due to the PLC's inability to process this request, it enters an error state, disrupting all communication to the HMI and actuators. No solutions were identified for preventing the PLC from entering an error state. Therefore, using Nmap's service and enumeration feature was deemed too high risk and should be categorised as prohibited during scoping of adversary-centric security testing for the ET-200S.

Similarly to the tests performed on the HMI in Section 6.2.2, running Nmap's service and version enumeration option, Nessus, and OpenVAS did not result in any abnormal behaviour; signifying that the use of these tools on the HMI does not cause Incompatible Data Basic Events.

6.3 Risk-Aware Scoping of ICS/OT Adversary-Centric Security Testing

6.3.1 Model Proposal for Zone and Level Scoping of Adversary-Centric Security Tests

While scoping of adversary-centric security testing for IT is often client-defined, the existing safety and operational risks discussed and quantified in Section 6.2, when conducting security tests within ICS/OT environments, provide further constraints for the scoping of these. This, therefore, requires further granularity to ensure that no disruption to the operational process is observed. As such, when defining the scope of an adversary-centric security test within

industrial environments, a layered methodology can be used to separate the scoping of zones and levels containing differing levels of risk.

To enable this, a hybrid model called the Testing in Depth for ICS (TiDICS) methodology is proposed, derived from the Purdue Enterprise Reference Architecture (PERA), also known as the Purdue Model, and the Defence in Depth Model. PERA is a commonly used architecture for segmenting devices and equipment within an ICS/OT environment into hierarchical functions. For the proposed framework model, an extended version of this which utilises a Demilitarized Zone (DMZ) to provide additional separation between the Enterprise and Manufacturing Zones has been selected [59]. Previously detailed in Chapter 1, its use can be applied to the scoping of adversary-centric security testing as the different zones and levels within the model, illustrated in Figure 1.1, have different risk levels due to the different device types implemented in each zone or level.

The second model used for developing the TiDICS methodology is the Defence in Depth (DiD) Model, illustrated in Figure 6.15. While the historical military strategy revolved around using weaker perimeter defence to allow the time to plan for a counter-attack, the cyber security strategy for DiD, conceived by the United States National Security Agency, involves parallel systems of physical, technical and administrative countermeasures to minimise the probability of a malicious actor gaining complete control of an environment [189]. Developed initially as a defensive strategy, the DiD model's layers can also be used as a testing methodology during adversary-centric security tests. These layers are as follows:

- Policies and Procedures: Cyber Threat Intelligence, Threat Modelling, Security Awareness Training, Security Governance, Risk Management, etc.
- Physical: Physical Access Control, CCTV, etc.
- Perimeter: Perimeter IDS/IPS/Firewall, DMZs, etc.
- Internal Network: Enterprise Remote Access, Content Filtering, Network Access Control, Data Loss Prevention, etc.
- Host: Patch Management, Endpoint Security Enforcement, Host IDS/IPS/Firewall, etc.
- Application: Database Monitoring, Dynamic/Static Application Testing, Application Firewall, etc.
- Data: Data Classification, Data Integrity Monitoring, Encryption, etc.

By combining these two models, The traditional testing methodology derived from the DiD model is now PERA Zone and Level dependant. For each layer of the DiD model,

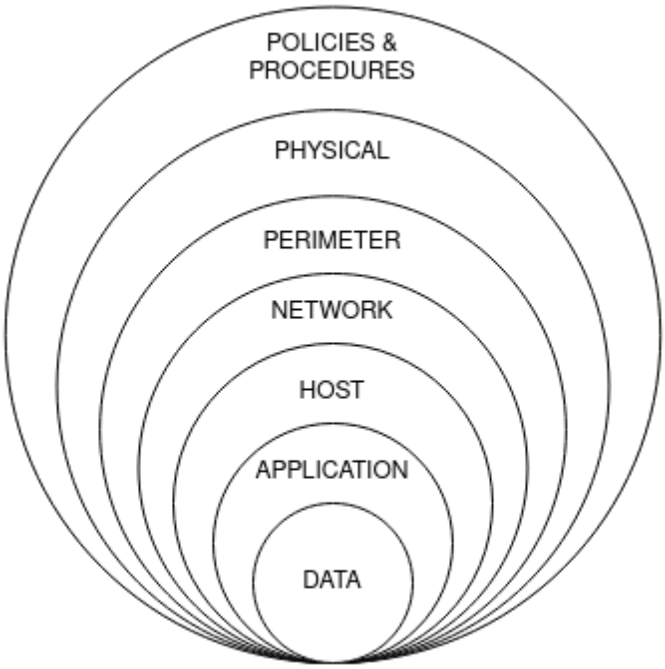

Fig. 6.15 Defence in Depth Model

the safety and operational risks within each zone or level from the Purdue model also need to be considered during scoping of an adversary-centric security test. With this additional separation of zones and levels, scoping can be done based on identified safety and operational risks, allowing for further depth of testing of zones and levels with fewer risk factors.

By adding PERA zone and level requirements to the traditional testing methodology for the DiD model, scoping of adversary-centric security tests can be granularised into separate testing levels with varying degrees of risk to the operational process. Therefore, tools and techniques used for testing can be defined on a per-zone and per-level basis, allowing for extensive depth of testing while ensuring that risk is minimised for each of these.

6.3.2 Framework for Risk-Based Scoping of ICS/OT Adversary-Centric Security Tests

By applying the methodology for identifying and assessing safety and operational risk of adversary-centric security testing, described in Section 6.1 and the TiDICS model for defining testing of zones and levels with various risk factors, the following risk-based adversary-centric security testing framework is proposed. The core output of this provides a methodology for integrating safety and operational risk into the scoping adversary-centric security tests within

ICS/OT environments; the extended framework (with example methodologies) is provided in Figure 6.17, and its process flow is provided in Figure 6.16.

The overall framework is used sequentially as the output of previous phases is used as input for subsequent phases. The following subsections provide a description of these phases, their input requirements and their outputs.

Select TiDICS layers

Depending on the type of adversary-centric security tests and the budget of the organisation being tested, relevant PERA zones and levels can be selected to facilitate scoping of these. Once these zones and levels have been selected, subsequent DiD layers can be selected for identifying and quantifying safety and operational hazards for each of these. For example, The entirety of the Cell/Area Zone can be selected for scoping of a security test. Following this, only network and host DiD layers are selected for testing. This signifies that an assessment of safety and operational risks for the following zones and layers needs to be undertaken to scope the engagement: Cell/Area network; Area Supervisory Control Network and Host; Basic Control Network and Host; and Process Network and Host. Cell/Area Host testing is arbitrarily removed from scoping as the scoping for level 0 to 3 host testing implicitly results in that of the overall zone.

Identify Safety and Operational Hazards

Identifying Safety and Operational Hazards that can be caused due to active adversary-centric security testing for each of the selected TiDICS layers needs to be undertaken next. Several methodologies exist for identifying risk events and can be used at the framework user's discretion. An example of identifying these risk events using a (Control) Hazard and Operational Study ((C)HAZOP) is provided in Section 6.1.1 and demonstrates how risk events can be identified using a guideword methodology. This phase is primarily qualitative and relies on existing documentation, such as P&I diagrams, network diagrams, configuration documents, and other relevant documents, for deducing risk events.

Decompose Safety and Operational Hazards

Once safety and operational hazards have been identified, these can be further decomposed into a combination of basic events that, if happen simultaneously, lead to a top event, or major hazard, occurring. Again, while the specific methodology for doing this is subject to the user's discretion, the framework provides an example Fault Tree Analysis methodology, discussed in Section 6.1.2. Conducting an FTA allows framework users to generate minimal

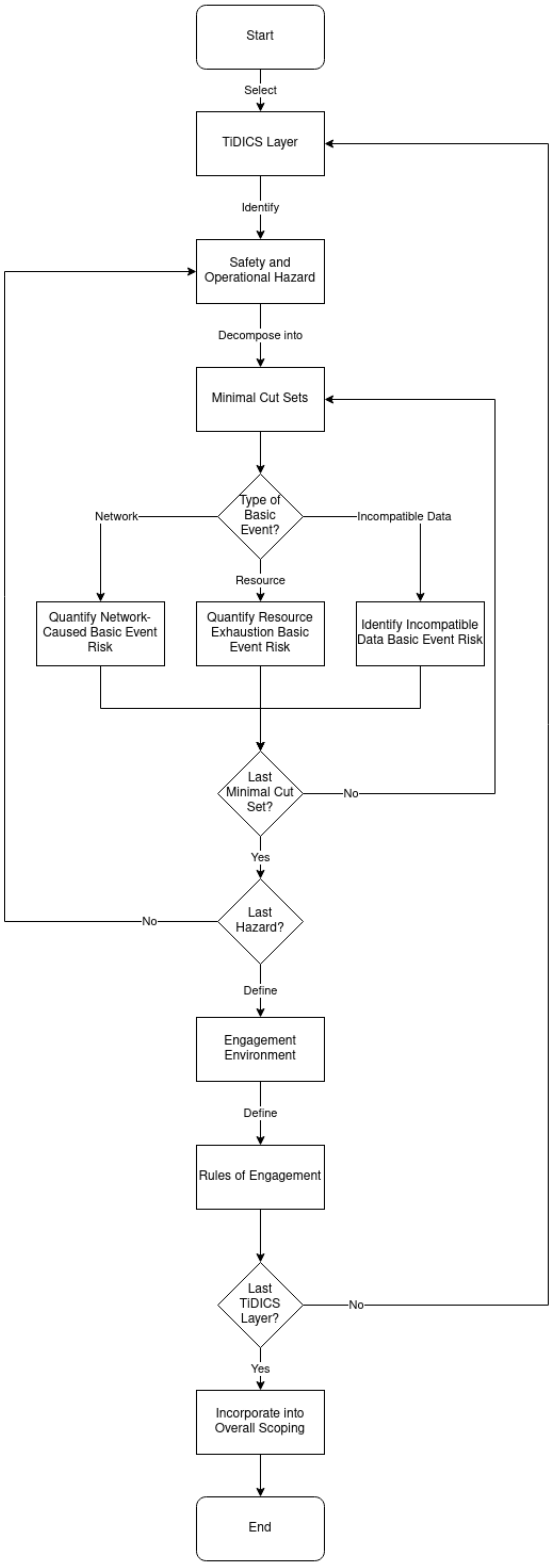

Fig. 6.16 Scoping Framework Process Flow

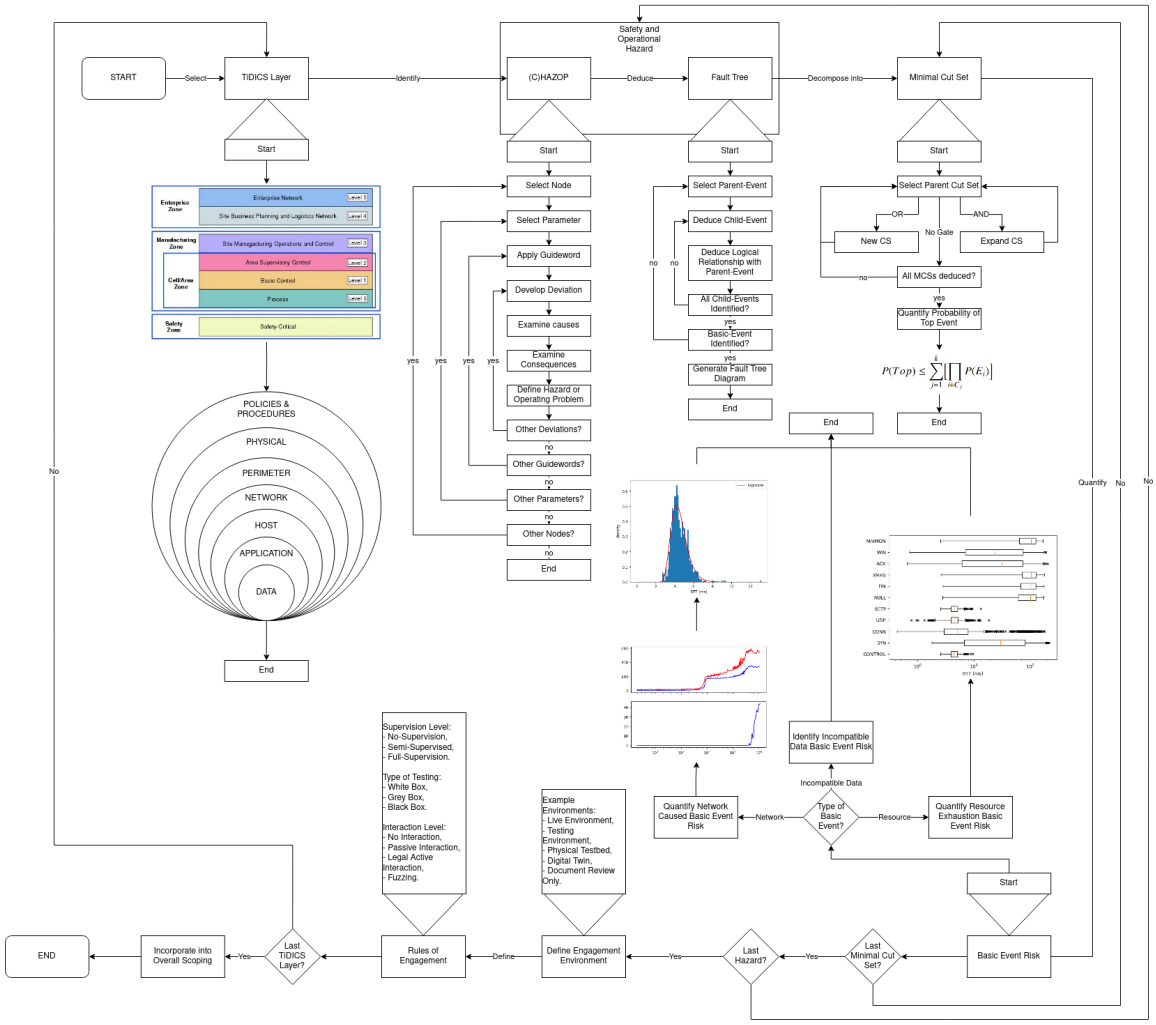

Fig. 6.17 Extended Framework for Safety-Risk-Based Scoping of Adversary-Centric Security Tests

cut sets of basic events, which can be represented through boolean algebra and therefore used to quantify risk precisely.

Quantify Network-Caused Basic Event Risk

Basic events that contribute to a reduction or loss in availability caused by latency increase or packet loss are categorised as Network-Caused Basic Events. The probability of these can be determined through prior testing or estimated by expert opinion. By determining the maximum allowed network throughput for specific endpoints or networks, the probability of tools or techniques affecting availability at these throughputs can be determined by considering network jitter. This probability of disruption to the operational process can be determined and used to determine appropriate security testing tools.

Quantify Resource Related Basic Event Risk

Basic events that contribute to a reduction or loss in availability caused by resource exhaustion are categorised as Resource Exhaustion Basic Events. The probability of these can be determined by testing tools and techniques planned to be used during the security test and determining their effect on the target. If tools have multiple options for performing a similar task, these can be compared to determine the options that produce the most negligible overhead and are the safest for use.

Identify Incompatible Data Basic Event Risk

Basic events that contribute to a loss of integrity or a total loss in availability caused by incompatible or anomalous data being sent to a target and unable to be processed or understood are incompatible Data Basic Events. Identifying risk mitigation techniques for specific tools or techniques that cause these basic events can be done through logical experimentation (i.e. changing tool options or improving exception handling on targets). However, if no solutions can be identified within an appropriate time frame, these tools and techniques must be documented, and their use during the security test should be prohibited to prevent disruption to the operational process.

Define Engagement Environment

Based on the quantified risks of selected adversary-centric security testing tools and techniques, the environment for deploying these can be selected. Such environments that can be used for testing include but are not limited to: live environment (high inherent risk and

high accuracy of testing results), testing environment (medium inherent risk and accuracy of testing results), physical testbed (low inherent risk and medium-low accuracy of testing results), digital twin (low inherent risk and low accuracy of testing results), and document review (no inherent risk and very-low accuracy of testing results). For example, if the risks quantified from previous steps are low, testing can be conducted either in a live environment or a test environment to ensure maximum depth of testing. However, if the risks quantified from previous steps are high, testing can be conducted in lower-risk environments such as testbeds to prevent disruption to the operational process.

Define Rules of Engagement

Based on the selected engagement environment, rules of engagement must be defined and subsequently enforced during the entirety of the security test. Several types of rules of engagement can be defined, such as supervision level (no-supervision, semi-supervision or continuous supervision), the type of testing (white box, grey box, or black box), the type of interaction with endpoints (no-interaction, passive interaction or active interaction) and more if required. These rules of engagement are categorised and defined within general scoping methodologies for adversary-centric security tests using comprehensive risk treatment principals, found in ISO 31000 [126], for example. These risk treatment options are not always mutually exclusive or appropriate in certain scenarios and can include risk avoidance, risk acceptance, risk removal, reducing likelihood, impact modification, or risk sharing. As described in ISO 31000, the use of risk treatment should take into consideration all party's obligations, voluntary commitments, views from stakeholders, objectives, risk criteria and available resources.

Incorporate into Overall Scoping Methodology

Finally, these results can be documented for use in the overall scoping of an adversary-centric security test to understand the safety and operation risks that are present. Incorporating the outputs of the framework in the overall scoping of the test ensures maximal depth-of-testing and minimal disruption to the operational process.

6.3.3 Discussion

The proposed framework provides a methodology for integrating the quantification of safety and operational risk within the overall scoping of adversary-centric security testing within ICS/OT environments. While the framework does not provide a specific methodology towards the integration of these risks, it offers a flexible approach to conducting safety and

operational risk assessment through example methodologies that can be used at the user's discretion. In addition, these example methodologies provide stakeholders with relevant direction to understand how their existing risk assessment practices can be used for the scoping of adversary-centric security tests using safety and operational risk quantification.

During the development of the framework, a few limitations were identified that could affect its effectiveness in practice. Firstly, the framework uses the TiDICS model, which incorporates both the Purdue and Defence in Depth Models. While the Purdue Model is highly recognised as a staple model for improving the overall security of an environment through zone separation, certain organisations' network architecture may not be fully modelled based on it; introducing additional complexity in the selection of zones and levels for safety and operational risk quantification. For environments that lack fundamental security principles, such as security-focused network architectures or a lack of asset management maturity, it is recommended that these be remediated prior to engaging in adversary-centric security testing engagements. However, for environments that use other reference architectures, which are often similar to the Purdue Model such as recommendations provided by IEC 62443 [119, 115], the TiDICS model can be modified to reflect these architectures instead.

Similarly, while these are provided as example methodologies for identifying and decomposing hazards within selected TiDICS zones and layers, the use of (C)HAZOP and FTA for safety and operational risk assessment can also present a few challenges during the implementation of the framework in practice. Firstly, (C)HAZOP is a qualitative risk assessment methodology by design. This means that to be able to quantify the risks that are identified using (C)HAZOP, FTA is also required. In order to effectively conduct risk identification and decomposition using these two methodologies, a minimal level of maturity in these areas is necessary, making use of efficient asset management and established safety risk assessment practices. If these maturity requirements are not met, the precision from the resulting output of (C)HAZOP and FTA may be insufficient, hindering subsequent steps of the scoping framework. Additionally, other existing methodologies may provide further precision than (C)HAZOP and FTA, depending on the target environment, such as a Systems Theoretic Process Analysis (STPA). Proposed by Leveson [155], this approach treats safety risks as control problems instead of failure problems which can provide further precision in control-heavy environments. Furthermore, STPA has been expanded extensively to include aspects such as security risk [80] or experiment specification [258], which can be used alongside the proposed framework to provide further depth in assessing the security capabilities of target environments.

6.4 Conclusion

The work described in Chapter 5 identified that the critical nature of Operational Technology environments presents several challenges for conducting adversary-centric security tests within these. Indeed, due to the inherent safety and operational risks within these environments and the design philosophy of systems such as Industrial Control Systems (ICS) deployed within these environments, simple tools and techniques that cause negligible effects within Information Technology environments could disrupt the operational process in industrial environments. However, recent technological advances in Operational Technology products such as the Siemens S7-1200 PLC [247], which have better resources than previous legacy products, have made adversary-centric security testing possible when considering the safety and operational risks that exist during these.

Within this Chapter is proposed a scoping framework for adversary-centric security testing, which incorporates identified safety and operational risk into the overall scoping process to minimise the risk of disrupting the operational process while maximising the depth of testing where possible. The first step in the design of this framework includes the proposal of a hybrid testing methodology, named the Testing in Depth for ICS (TiDICS) model, created through the combination of the Purdue Enterprise Reference Architecture [59], also known as the Purdue Model, and the Defence in Depth Model, conceived by the United States National Security Agency [189]. While the Defence in Depth model is traditionally used to implement several layers of security controls, this can inversely be used as a methodology for testing separate layers of an environment. The Purdue Model is a reference architecture for the integration of Operational Technology into a broader enterprise network. By segmenting devices into hierarchical functions, levels can be defined with different safety and operational risk factors. By combining these two models, a testing methodology can be devised for conducting adversary-centric security tests on a per-layer basis with unique risk considerations for each level within the Purdue Model.

After defining the layers for testing, identifying and assessing the safety and operational risk within these can be done. Several methodologies exist for identifying operational hazards, and an example methodology using a (Control) Hazard and Operability ((C)HAZOP) study is provided due to its widespread application across several domains, including Critical National Infrastructure [65]. Once hazards have been identified, these can be further decomposed into basic events, which, through Fault Tree Analysis, are a set of events that, if occurring simultaneously, lead to a hazard.

Three types of basic events were identified from conducting adversary-centric security testing: Network-Caused Basic Events, Resource Exhaustion Basic Events, and Incompatible Data Basic Events. Network-Caused Basic events occur from employing tools and techniques

that produce excessive amounts of network traffic, which lead to either congestion of the network or an increased response time from endpoints being tested. This results in a reduction in availability and, therefore, could disrupt the operational process. Resource Exhaustion Basic Events occur from employing tools and techniques that produce data that cause additional overhead on an endpoint's resources, leading to an inability to operate on time. Finally, Incompatible Data Basic Events occur from employing tools and techniques that produce data that cannot be processed by a target endpoint, resulting in it entering an error state or crashing and resulting in a total loss of availability. Quantifying the risk of these basic events can be done through testing and interpolation techniques within a physical testbed.

The identified and quantified safety and operational risks for each layer of the TiDICS methodology are then used to define the engagement environment and the rules of engagement for these, including scoping constraints. These can then be incorporated into the overall scoping of an adversary-centric security test so that safety and operational risk is minimised and depth of testing is maximised.

While there are limitations with the usage of the proposed framework, these are primarily due to the possible lack of safety and operational risk analysis maturity from target users. To this end, an evaluation of the implementation of the framework in practice is required to assess its accuracy, reliability, validity, and applicability based on existing processes. As part of this evaluation, presented in Chapter 7, semi-structured interviews with key stakeholders, including ICS/OT cyber security consultants, engineers and penetration testers, were conducted; these roles being directly associated with the implementation of the framework.

Chapter 7

Evaluation

Chapter 6 introduced a framework to aid stakeholders in scoping adversary-centric security tests within ICS/OT by quantifying safety and operational risk. The framework provides a step-by-step methodology for identifying safety and operational hazards, decomposing these hazards into risks, quantifying these risks, and defining scoping constraints for integration into the overall scoping methodology of these engagements. While the framework is supported by commonly adopted methodologies and data collected from the Lancaster University ICS testbed, these only provide a proof of concept of how the framework could be adopted and implemented in practice. Therefore, this chapter provides an evaluative study of the framework to determine its potential for implementation in practice.

While current literature and practice acknowledge that safety and operational risk are a concern when conducting adversary-centric security tests in ICS/OT environments, these do not address how to identify and understand these risks to reduce the impact of these engagements on safety and the operational process. The risk-based safety scoping framework was created to address this gap to enable safe adversary-centric security tests within ICS/OT environments. For the evaluation of the framework, a qualitative study with participants involved in ICS/OT adversary-centric security testing was selected. This allowed for collecting data on experts' opinions on the accuracy, reliability, validity, and applicability of the framework's implementation in practice while also identifying potential issues that could limit its use.

7.1 Methodology

Due to the critical nature of the context in which the scoping framework is intended for, a high level of confidence is required before it can be tested and evaluated in real industrial environments. To this end, a qualitative approach was selected in which interviewing

stakeholders across the topic area of ICS/OT adversary-centric security testing was applied, similar to the chosen methodology used to assess the implementation of standards and guidelines for cyber incident R&R in practice in Chapter 3. For this, participants were provided with the framework and an example application of its use, using data collected from a testbed environment in the context of the water tank scenario described in Chapter 6. This approach provides two benefits: firstly, participants were able to provide their opinion on the framework, its phases, as well as any observations on potential limitations that could hinder its use in practice; and secondly, using the data collected from the aforementioned scenario, participants could provide insight on the framework's accuracy, reliability, validity, and applicability for implementation in practice. These evaluation criteria have been defined as follows:

- Accuracy: How close is the output of the framework to the correct and accepted outcome?
- Reliability: How does repeated use of the framework affect its outcome when provided the same input?
- Validity: How appropriate is the framework in addressing its objective?
- Applicability: How much does the outcome of the framework change when used in different contexts?

Interviewing was selected as an appropriate method that enables each participant to discuss their personal experience concerning the risks of adversary-centric security testing in ICS/OT environments and how the framework can be used to address these concerns [216]. For this, a semi-structured approach was also adopted as it provides adequate flexibility with a core question set while allowing the option to include improvised follow-up questions for further exploration of topics of interest [13]. Furthermore, the threats to validity concerning the reliability of the collected data using this method have previously been addressed within Chapter 3 Section 3.4.1.

7.1.1 Sample

The aim of selecting an appropriate participant sample is to understand the topic area from all relevant perspectives. To achieve this, a broad approach was applied to target participants. This resulted in a diverse collection of role profiles. More specifically, roles that would engage in the scoping or implementation of adversary-centric security tests within ICS/OT environments across multiple backgrounds and with varying levels of responsibilities. This

sampling approach provides multiple perspectives, building an accurate perspective on the scoping framework's validity for implementation in practice.

To summarise, five participants were selected holding the following roles:

- ICS/OT Security Researcher
- ICS/OT Cyber Security Engineer
- Health and Safety Manager
- Operations and Finance Chief Information Security Officer
- Filling and Packing Operations and Cyber Security Manager

The levels of experience varied amongst participants within each of the defined roles ranging from five to forty-six years; the majority of which, however, had been working with industrial systems for over ten years.

While having a sample size of five for this evaluation may seem insufficient for the accurate evaluation of the framework, the main purpose of this evaluation is to identify limitations with its implementation in practice. Nielsen et al. provide a mathematical model for finding usability problems that can be used to plan the amount of evaluation required to achieve the desired level of thoroughness [200]. This work demonstrates that by conducting a qualitative study with five participants, 85% of the issues in a proposed work, such as the scoping framework, in this case, will be identified. Ensuring that a 31% chance exists that each participant will identify an issue if it exists, is also defined as a requirement. Given the role profiles and the levels of experience of the selected participants for the study, this requirement is met as each person has sufficient expertise to identify any issues with the framework. Therefore, having a sample size of five for the evaluation of the framework is appropriate for identifying at least 85% of the possible limitations concerning the implementation of the framework in practice.

7.1.2 Interview Protocol/Guide

Each interview was broken down into the following seven stages, providing a logical structure to the interview protocol/guide:

- Preface
- Establishing Demographics
- Framework Familiarisation

- Framework Evaluation
- Scenario Familiarisation
- Framework Evaluation (With Application Scenario)
- Conclusion

The core focus of these interviews was to present participants with the risk-based safety scoping framework for adversary-centric security testing on ICS/OT. More specifically, how key stakeholders would approach risk quantification for the scoping of these engagements using the framework. Taking direction from the phases of the framework, discussed in Chapter 6, the questions aligned to these interview stages are aided through the inclusion of probes and definitions. The following provides a summary of primary interview questions with the complete protocol/guide provided in Appendix B.

Establishing Demographics

The following question-set was applied to the demographics phase:

- Please can you tell us your job title and provide a brief overview of your core roles and responsibilities?
- How many years experience do you have working in this role?
- At a very high level, please can you explain what you understand the term adversary-centric security test to mean?
- Have you ever been involved in an adversary-centric security test that was performed for an ICS/OT environment?
- At a very high level, what do you believe to be the greatest challenges of conducting adversary-centric security tests for ICS/OT environments?

Framework Evaluation

The following question-set was applied during the Framework Evaluation phase once participants had been presented with the risk-based scoping framework:

- What is your opinion on using the TiDICS model for separation of testing zones and layers based on safety and operational risks?

- Do you agree that the types of risk that adversary-centric security testing presents to safety and the operational process are comprehensively considered within the framework?
- What challenges could affect the collection or quality of data for risk quantification?
- From the framework's overview, do you think the output of the framework could be used in the overall scoping of an adversary-centric security test?
- From the framework's overview, do you believe that the output of the framework is accurate enough to ensure a full understanding of the safety and operational risks from adversary-centric security testing on ICS/OT so that depth of testing can be maximised while minimising risk to the operational process?
- From the framework's overview, do you believe that the framework can be applied in all ICS/OT environments where safety and operational risks are a concern?

Framework Evaluation (With Application Scenario)

The following question-set was applied during the Framework Evaluation (With Application Scenario) phase once participants had been presented with an example scenario for applying the framework:

- Does your opinion of the framework's accuracy, reliability, validity, and applicability change when presented with an example application of its use?
- What is your opinion on the use of (C)HAZOP for identifying hazards that could occur during an adversary-centric security test on ICS/OT?
- What is your opinion on the use of FTA to decompose hazards into smaller basic events for use in risk quantification?
- What is your opinion on the methodologies used for quantifying the risk of basic events?
- Overall, would you use this framework as part of the overall scoping methodology for an adversary-centric security test on ICS/OT?

Conclusion

The following question-set was applied during the conclusion phase:

- Would you like to add anything which may be relevant?

7.1.3 Analysis

Similar to the methodology selected for analysis of the interview data from the stakeholder engagement on cyber incident R&R in Chapter 3, template analysis was selected. To summarise: Also referred to as "codebook analysis" or "thematic coding", template analysis is used as a highly flexible method for analysing qualitative data, which was collected during these interviews [145]. This approach is considered a middle ground between the relatively rigid content analysis approach in which analytical codes are all pre-defined [229] and the opposite approach of grounded theory in which all analytical codes must be derived from the data [88]. Through this methodology, an initial code set can be created based on the interview protocol/guide, which is aligned with the core areas of interest for evaluating the framework. Furthermore, template analysis allows for creating additional code sets to analyse discussion points previously not considered.

7.2 Results

Key findings from the interviews with stakeholders are provided here. Similar to the interview results from Chapter 3, Section 3.4, these have been grouped based on identified themes and key points of interesting. All of the points discussed here have been selected based on the contents of the interviews. These, therefore, reflect the generalised opinions of participants but may differ from person to person.

7.2.1 Challenges of Adversary-Centric Security Testing on ICS/OT

Before evaluating the framework, it is important to understand what participants understand to be the current challenges of conducting adversary-centric security tests within ICS/OT environments. Participants' opinions aligned closely with the findings from Chapter 5. Due to the nature of ICS/OT environments, any engagement that can potentially affect the operational process can lead to a loss of business continuity or even safety.

'When you are operating an OT environment, everything needs to be run as efficiently and effectively as possible and any slight deviations from that can not only put human safety at risk, but it can also completely throw off the operational process and in turn cost a lot of money.'

The safety and operational risks are amplified due to the design philosophies of ICS, as discussed in Chapter 2. Generally, these are designed to favour environmental resilience and operational longevity instead of performance. These limited resources could affect

availability in the context of an adversary-centric security test if using tools that are resource heavy.

'They have very limited resources, their design tends to be relatively outdated, and then it leads to the fact that a lot of these devices are used in critical environments.'

The timing for conducting adversary-centric security tests within ICS/OT environments was also discussed. Due to the high up-time requirements of these environments, maintenance periods for live environments are uncommon, which can occur every five to ten years in some environments. Because of this, a decision has to be made on whether engagements can and should be conducted outside maintenance periods or should be restricted to these.

'You'd have to wait until plant shut down there. They'll have routine maintenance which can be every six months but it could also be every five to ten years'

Due to ICS/OT cyber security, in general, being an interdisciplinary field, several different skill sets are required for successfully conducting adversary-centric security tests. Therefore, ensuring that these different actors are present was identified as crucial by participants to ensure all required considerations during the planning and execution of these engagements.

'Generally, There are three actors that don't all have the same skill set. The IT team in general doesn't really know the OT world but do know cyber security. The engineers know how the systems are used but the people that know exactly how the devices work are the vendors.'

This is equally applicable to the test providers, which also require ICS/OT knowledge in order to successfully understand how to provide the correct services for these environments.

'When conducting a pen-test on ICS/OT, the testers need to know what an industrial device is since they use different protocols and are designed differently.'

Because of the wide variety of protocols and product vendors used for ICS/OT, difficulties also arise when conducting adversary-centric security testing within environments that incorporate a broad range of protocols and device types. This adds additional complexity to these engagements due to different protocols and devices often being incompatible with each other. For example, if Modbus and S7COMM were used within the same network, this could require different tools for conducting tests which leverage these protocols, such as SimaticScan [11] for S7COMM.

'Because you have different types of devices, if you have a unique way to test these with one tool, that would be helpful for the OT team and allow for comparable results between devices.'

7.2.2 Selection of Testing Zones

Despite the Purdue Model and DiD Model not being initially intended for scoping adversary-centric security tests, participants agreed that these could be used for defining testing areas for safety and operational risk analysis. While used as a reference architecture, the Purdue Model separates zones and levels based on hierarchical function. Because of this, the different zones and levels that can be selected will also contain different risks, enabling distinct risk analysis processes for each of these.

'I think it provides you the capability of at least defining which devices you're going to test. And definitely the further down [the Purdue Model levels] the more things are critical.'

As discussed in Chapter 6 Section 6.3.3, while the Purdue Model is commonly adopted for designing ICS/OT networks, other reference architectures are also used in practice, such as the one recommended by IEC 62443 [118]. For this, the framework can be modified to select testing zones based on these reference architectures, allowing flexibility for environments that use different network architectures than the one described in the Purdue Model.

'I would definitely look at how IEC 62443 handles zones and conduits because I think it would make the framework even more accessible for environments that are already following 62443.'

Similarly, the DiD Model, while used for implementing different defensive controls and policies on a per-layer basis, can also be used to further granularise the selection of areas to test. The risk for each of these layers can then be analysed for use in subsequent phases of the framework. A few participants noted the dependency between certain layers of the DiD model depending on the design philosophies of the networks or devices being tested. For example, when testing the application layer of a target device, this might imply that testing at a host level overall is also required due to the architecture of the target device(s).

'[The DiD Model] might be a bit more difficult on OT because sometimes application and data meshes into the host itself if it's an embedded system. You'll have things like Real Time Operating Systems, for example.'

The TiDICS Model could facilitate risk identification and analysis for different environments with a similar configuration. This would allow framework users to streamline specific steps of the framework for multiple environments, reducing the cost of the scoping process for adversary-centric security tests.

'It allows us to deep dive into the risk of a specific system and replicate that somewhere else with similar configurations. It allows us to do a safety risk assessment that's reproducible in similar environments.'

7.2.3 Framework Users

While the framework was initially intended to be used by both the test providers and the environment asset owners (including IT and OT engineers), a third party was identified that would need to be involved in the scoping process to maximise its efficiency: the product/solutions vendor. This is because vendors provide certain environments as part of a black-box solution.

'These live environments are often black box in the sense that the solutions for this are provided by the vendor.'

Because of this, certain organisations' engineering teams may understand how their environments function but would need to gain the required knowledge concerning the inner workings of specific devices for in-depth risk analysis.

'There'll be engineers that understand what the OT devices are doing but how they're built inside such as the OS, the protocols, and the connection between other devices; only the vendor knows because they're the ones that provide us with a solution that's ready to use. [...] The engineers only understand the functional aspect of their environment'

This lack of required information would also include crucial documents used for risk identification, such as network diagrams, which are essential for comprehensively considering the risk that conducting an adversary-centric security test could present to these environments.

'For example, the network diagrams - we don't have that. It's the vendor that has that. They could give us that information if required but often it's them that understands fully how the PLCs and networks are configured.'

7.2.4 Safety and Operational Hazard Identification

As part of the initial phase of the framework for safety and operational risk identification, difficulties could arise for organisations without proper asset management maturity. This would impact the quality of hazard identification and could lead to significant hazards not being identified.

'if we actually know what we have in the network. [...] You need to know what you've got in order to be able to analyse it.'

Identification of hazards would also need to include the product/solutions vendor to fully understand the risk when conducting an adversary-centric security test.

'The engineer can identify the risks of systems during operation but not on the way that they're configured. So for identification of risk, the vendor would need to be involved.'

While only a few participants were aware of (C)HAZOP prior to the demonstration of the framework, its ease of use and high-level operation allow it to be used within several contexts, including adversary-centric security testing. Applying guidewords to parameters to identify potential deviations allows framework users to comprehensively consider all risks that could be present during such an engagement.

'It's interesting in the sense that we take each parameter and apply a guideword to identify hazards. [...] It's more precise than say "if the HMI doesn't work anymore, what happens". Because CHAZOP provides such a high level methodology, it can be applied within a lot of different contexts, which is why it's good to use. '

Participants appreciated that the framework offered flexibility on the methodologies for identifying hazards, enabling the use of methodologies already used for traditional safety risk assessment in the context of adversary-centric security testing. Using hazard identification methodologies that are well established and already used by the framework users also facilitates this phase of the framework since a new methodology would not need to be learnt from the beginning.

'I think it's always best to stay with the things that are actually used in practice and have reputation. So I can also imagine that you would be able to use this framework with other methodologies as well. [...] In the chemical industry, we also use something called bow tie analysis.'

For organisations that do not have an established methodology for identifying hazards, providing (C)HAZOP as an example methodologies offers framework users a starting point in the event that more guidance is required for this phase of the framework.

'Our preference is HAZOP. There are some other methodologies, but I think what's important is that there is a methodology out there that can be used for pen-tests. It's a good methodology to use but what's important is that we identify a methodology that can be used rather than just trying to do something blindly.'

7.2.5 Risk Initiator Deduction

To enable risk quantification in subsequent phases of the framework, identified hazards need to be broken down into smaller quantifiable risks. While not all participants were familiar

with Fault Tree Analysis, most participants were familiar with methodologies that use tree diagrams, such as attack trees or probability trees.

'I think that most of the people that would be involved would know how to read a tree representation. It's fairly common across multiple disciplines.'

FTA was deemed adequate for decomposing hazards into smaller risks as it provides a logic tree that describes relations and dependencies between different risks. Being environment-agnostic, it is also widely applicable to several contexts, including adversary-centric security testing.

'FTA is a powerful tool to take into accounts all the devices, steps, causes and consequences. So for me, it's adaptable and the power of that is that it's not linked to any specific environment. You can use it for areas like filling and packing, manufacturing, flows, etc.'

Because FTA uses logic tree diagrams, these can be directly translated to probability tree diagrams. Subsequently, these can be used to calculate the probability of identified hazards occurring based on smaller quantifiable risks.

'You'll be able to determine the probability of something going wrong in a live test. So having that kind of information for me gives a comfort factor during a penetration test. Even though you might not be able to tell me it's 100% safe, it could be 97% for example.'

Similar to how (C)HAZOP is provided as an example methodology for identifying risks, decomposing hazards can be done using other methodologies if they achieve the same result of being able to quantify risks.

'There are a few other methodologies out there like failure mode analysis. [...] If a company worked on a different type of system to achieve the same result, you could accept that.'

However, FTA was accepted as a suitable recommended methodology and should be used if there are no established methodologies for decomposing risk.

'On the other hand, if what you're trying to do is put out a methodology into a new space, then it's worthwhile saying "OK, this is the recommended way to do it" when people don't know what they're choosing. Sometimes it's better to say "this is what we recommend". [...] FTA is a simple one and it's typically very easy to understand.'

One concern that was raised was the scalability of using these methodologies for adversary-centric security tests with large scopes. While the scenario provided was simple in concept, some industrial environments can be very large and complex, which could lead, in turn to very large and complex fault trees being generated. In this case, the framework does acknowledge scalability issues and suggests mitigating these by using fault tree generation software to automatically create FTDs.

'Depending on the size of the scope, this sort of thing could become very complicated very quickly. [...] I think it could become very time and resource demanding to do this.'

7.2.6 Collecting Data

In order to efficiently quantify the risks of conducting adversary-centric security tests within ICS/OT environments, the appropriate data needs to be collected. However, the skill gap between different fields of expertise could lead to vital information being missed for risk quantification. This highlights the importance of including all relevant parties in the scoping process so that the collected data comprehensively covers all the risks present during these engagements.

'It's very difficult to find somebody who's an expert in OT, risk and cyber security at the same time. So for me one of the main issues is generally the knowledge of the people involved in the scoping process. Who needs to be involved so that it [the collection of data] can be done to an in depth extent?'

While passively collecting data in live environments is possible, some environments might be provided as black-box solutions and require the product/solutions vendor to provide appropriate data or install solutions for passively collecting data, such as maintenance sensors. In some instances, however, the vendor may be unwilling to implement these.

'We need sensors that passively analyse the network. For these, some vendors have partnerships with companies that provide these services but other vendor might not allow that. And others might say to not touch anything because then they would no longer guarantee that their provided solution works as advertised.'

In order to safely collect data for quantification of risk through active methods, doing so through a staging environment or testbed was recommended in order to prevent potential disruption within a live environment.

'It's better to do it offline in a testbed environment than actually out on the shop floor. You can take it, isolate it, test it and find out whether there's any problems before you actually start testing on a live environment where it could have a catastrophic effect.'

The quality of collected data from an isolated environment could impact the precision of the quantified risks. While this is true, it was generally accepted that if the offline environment (such as a testbed) closely mirrors the live environment, the margin of error in risk quantification would be negligible.

'I don't think there's no chance that it could act differently to a PLC that is in an environment under load with additional network traffic. So there are some situations where, if I was extremely risk averse, I might be concerned about this. But you could probably get as close as realistically possible by utilising the methods that have been described.'

However, this margin of error reduces dramatically for risks with high impacts. If a PLC crashes or hangs during an offline test, for example, it is likely to react the same within the live environment.

'If on the testbed, the testing shut something down completely 100% of the time, then it would have the same effect on the shop floor. [...] For more intricate events that might not cause that much impact then that might affect the quality of collected data. But for major events, then you could directly correlate between the two environments.'

Identifying the appropriate data to collect is also important. With modern technology, however, the amount and quality of the collected data could cause delays during this phase of the framework. Therefore, understanding which data is vital for risk quantification is crucial to prevent unnecessary complications.

'What data would actually be useful for evaluating risk? How do we collect this? How would it be used to perform an analysis? If it's an old system, do we actually have the capability to collect data for this? [...] Since ICS/OT, when implemented, is designed to stay operational for years, that could make things difficult.'

7.2.7 Methods for Risk Quantification

When discussing example methodologies for quantifying the risk of basic events, described in Chapter 6, the results from these would allow testers to determine precisely how tools and techniques could affect safety or the operational process. In doing so, additional information could be identified, such as the effect of network-based attacks on an environment.

'This allows us to determine which tools have a tendency to have a high throughput and which ones would be useable for the devices being tested. The actual hacker doesn't care. If they're wanting to cause disruption, the more disruption the tool causes the better it is for them. But if it's for a pen-test, the objective of that would be to verify vulnerabilities without affecting latency.'

The depth of analysis provided by the example methodologies lead a few participants to question their current process for scoping adversary-centric security tests within ICS/OT environments. By conducting a more in-depth analysis with quantifiable results, risk can be precisely calculated and used in the scoping process, improving the depth of testing while minimising risk to safety and the operational process.

'When we do this kind of test, we don't have this deep of an analysis. Maybe that means we take high risk. This is something that we would need to perform to identify areas we need to avoid during testing.'

Mapping the initial objectives of the adversary-centric security test to the methods used for risk quantification was also identified as vital. In doing so, associated risks can be minimised to prevent disruption to business continuity risk or vendor maintenance risk, and testing quality can be improved.

'This allows us to use the correct tools for the given scope but we would need to associate them with the initial objectives that were defined. For me, there's two things that we need to reduce: the risk associated with business continuity and vendor maintenance risk. If the vendor says: "your pen-test has changed everything based on our initial configuration and that means we can't ensure proper maintenance of the network" then that's very bad.'

While some scalability issues were identified for environments containing a wide variety of protocols and device types, participants generally agreed that the methodologies provided in the framework were a good starting point for quantifying safety and operational risk and should be used for the scoping of adversary-centric security tests on ICS/OT environments.

'Even if you've got 30, 40 or even 50 really important processes that have different systems, then I would say you need to go through them all using this approach and say "OK, what is the impact? What can I do? What should I be doing?" [...] Even though it could potentially be a lot of work, if you don't do it, you don't know what risks you have.'

7.2.8 Framework Outputs

Once safety and operational risks have been quantified, these can be used to define scoping constraints for use in the overall scoping methodology of an adversary-centric security test on ICS/OT environments. Such constraints include defining the engagement environment, for example. While testing on a live environment is possible, providing that the risk is sufficiently low, a few participants stated that despite this, some stakeholders may still be too risk-averse to allow any sort of engagement.

'I'm not sure how far people want to do tests on the live environment just yet. Although it might be more of a possibility in the future given more open-minded people and more modern technology.'

On the contrary, some participants stated that they would only be able to conduct tests in a live environment due to the lack of a staging area or testbed. Again, this emphasises the need to quantify safety and operational risks so that scoping constraints can enable safe adversary-centric security tests within live environments.

'We don't have a testing or staging environment. We don't have any environment other than the live environment. So either we ask to do some testing in an equivalent environment that the vendor has, if that exists, or we have to do it in the live environment.'

For this reason, another constraint identified as essential for scoping adversary-centric security tests is when to conduct the engagement. While specific environments may have long timeframes between maintenance periods, these can be used to conduct security tests with lower risk to the operational process than outside of these periods.

'What's important for me is when we're going to do the pen-test. This is very important as well because ICS/OT never stops. Sometimes, we have maintenance periods and we could do that then.'

Since risk is quantified in previous phases of the framework, precise scoping constraints can be defined, such as limiting network throughput on specific tools. This allows framework users to make accurate decisions that can be re-evaluated post-engagement for future security tests.

'It's very useful to be sure that we address all the different steps and take into account all the outputs. And that can help to compare the results of risk quantification to the expected results after the test.'

7.2.9 Framework Discussion

Overall, participants were mostly receptive to the framework and its use in the risk-based safety scoping of adversary-centric security tests on ICS/OT environments. However, a few limitations were identified with this. Firstly, the maturity of framework users may impact its effectiveness. Because of the critical nature of ICS/OT environments, stakeholders may be unwilling to conduct security tests within their environments even if the quantified risk is low. However, with modern technologies and methodologies for quantifying these risks, the framework presents an opportunity for safely conducting adversary-centric security tests within ICS/OT environments.

'It's not a limitation of the framework itself, it's a limitation on the maturity of the users of the framework. I probably would say that most people would be a little too cautious than actually they should be. The more people use the framework and the more knowledgeable they become about this sort of thing, then the more people might be willing to take risks of doing these tests in these environments.'

As discussed in previous sections, scalability was also identified as a potential framework limitation. The resources and time required for effective scoping of adversary-centric security tests may be too costly for large and complex environments. However, the framework provides users with a starting point for quantifying safety and operational risks to enable security tests in these engagements, especially for organisations that require such tests as part of compliance requirements.

'I think the whole battle of this framework would be timeliness of actually performing the framework for the diversity of the TTPs and assets versus the accuracy of the output. Despite this, I think a lot of organisations would like to have penetration tests. No one's done something like this framework and I think it's something that the industry is crying out for. It would be something that would be readily accepted because of the difficulties that this would begin to overcome.'

While the framework is intended initially to aid in scoping adversary-centric security tests for ICS/OT environments, it also serves as a tool to increase an organisation's asset management maturity. This can, therefore, be used in other areas, such as general safety risk management or improving operational asset resilience.

'It's not just enabling a penetration test, it's discussing and finding out information about the assets within an environment. Not only would clients be provided with a penetration test at the end of this, they'd also know more about their assets and it would enrich their asset register.'

Despite the framework requiring technical competence for its operation, participants appreciated that it was presented in a way that could be understood at all skill levels. This allows the framework to include both technical and non-technical users in scoping adversary-centric security tests within ICS/OT environments.

'It's simple and I don't mean that as in it's easy. You've managed to make a very complex process become simple. It seems to be adequate to explain the scoping to both engineers and the security operations centre for example - and even senior (management) people. It would probably put their mind at rest.'

While the framework provides flexibility in the methodologies used for risk identification and quantification, it offers recommended methodologies, including (C)HAZOP and FTA, that can be used following the provided guidance.

'With HAZOP and FTA you can almost plug and play the framework. These could maybe be changed to something else if needed but I think as a methodology, what's important is to put something out there in an area where I haven't really seen much.'

The framework's efficiency can also be improved the more it is utilised through lessons learnt. For example, assessing whether the defined objectives prior to scoping the adversary-centric security test have been met can be used to improve upon future engagements and the quality of risk scoping.

'We define certain objectives, we do the pen-test and then after we see if we met these objectives. And that's important because it allows us to improve the next pen-test and to improve the quality of risk scoping. A lessons learnt to see if the framework managed to appropriately identify the risks or if there are any gaps and how to resolve these.'

Finally, most participants shared that, having been presented with the framework, they would insist that such a process must be used prior to conducting adversary-centric security tests within their ICS/OT environments.

'I would actually insist that something like this was done, if it could be done. I think that it could be a requirement during the call of offers from pen-test providers to ask that they can ensure that the tools they are going to use aren't going to affect business continuity.'

7.3 Conclusion

Chapter 4 introduced the risk-based safety scoping framework for adversary-centric security tests within ICS/OT environments. This framework is designed to aid stakeholders in identifying and quantifying safety and operational risks so that adversary-centric security tests on ICS/OT can provide maximal depth of testing while preventing impact on the operational process and safety. Across this chapter, the methodology applied to interviews with participants involved in ICS/OT adversary-centric security testing has been outlined. This included a pre-defined question set, allowing for a degree of flexibility through semi-structured interviews. During these interviews, participants were presented with the framework and an example application of it using a fictional scenario designed in the Lancaster University ICS Testbed, described in Chapter 6. The intended goal was to assess the framework's accuracy, reliability, validity and applicability for implementation in practice as well as identify any potential limitations that could hinder this. The output of these interviews was analysed using template analysis, which was deemed a suitable technique given the nature of the defined research objectives.

Several themes emerged from this analysis, covering the phases defined within the risk-based safety scoping framework for adversary-centric security tests on ICS/OT and other relevant topics, discussed in Section 7.2. Overall, participants found that the phases presented in the framework were appropriate and realistic for use in practice and addressed the existing challenges of conducting adversary-centric security tests within ICS/OT environments. The selection of testing zones and layers through the TiDICS model, which leverages both the Purdue and DiD models, allowed participants to identify and separate areas to test, enabling further granularity when identifying and quantifying safety and operational risks. The model's flexibility to include other reference architectures, such as the one provided by IEC 62443 [118], was highlighted as providing additional benefits to framework users. An important topic which was initially overlooked in the current iteration of the framework was the inclusion of product/solutions vendors in the scoping process. Because ICS/OT environments are often provided as black-box solutions in practice, including vendors in the scoping process would add further depth to identifying and quantifying risk and improve the quality of the framework's output significantly. For hazard identification and risk deduction, participants appreciated the flexibility that framework users can use methodologies already implemented within the organisation. In addition, the inclusion of (C)HAZOP and FTA as recommended methodologies was also valued as it provides framework users with guidance where required. The example methodologies for collecting and using this data to perform risk quantification were recognised as providing value to the framework's reliability. From this, precise scoping constraints can be defined, allowing for adversary-centric security tests that

are scoped using this framework to ensure maximum depth of testing while minimising risk to safety and the operational process. One participant stated that the framework's validity and application go beyond the scoping of adversary-centric security testing as it can also enrich an organisation's asset management process, providing further use than initially intended.

While there was an overall positive response to the framework, a few concerns regarding its applicability in practice were identified. Firstly, the maturity of the users of the framework was highlighted. If implemented incorrectly due to a lack of maturity from the users, the quality of the framework's output could be less than intended. While this is not a limitation of the framework itself but rather a limitation of the framework users, this is acknowledged and mitigated through provided guidance for each phase of the framework. The main concern that could affect the framework's applicability is its scalability for large, complex and varied environments. This limitation has been acknowledged, and solutions for this will be considered in future work, discussed in Chapter 8. Despite this, participants agreed that even for large and complex environments, the framework provides a starting point for organisations to include safety and operational risk quantification in scoping adversary-centric security tests on ICS/OT.

In summary, the evaluation was considered a success, with results from the analysis of the interviews aligning with initial expectations of the framework's development. The risk-based safety scoping framework for adversary-centric security testing on ICS/OT was determined to be sufficiently accurate, reliable, valid, and applicable to provide a significant benefit to the overall scoping of these engagements.

Chapter 8

Conclusion and Future Work

In Chapter 6, a risk-based safety scoping framework for adversary-centric security tests on ICS/OT was introduced. This framework aims to address the gaps identified in current literature and practice where adversary-centric security testing for improving cyber incident R&R is uncommonly employed within ICS/OT due to the safety and operational risks that exist within these environments. The framework's utility was supported using data collected in an ICS/OT testbed composed of real-world hardware used in industrial environments. However, it was felt that further evaluation should be carried out to help assure the framework's validity.

Chapter 7 presented an evaluative study in the form of semi-structured interviews with experts in both ICS/OT and adversary-centric security testing. In this qualitative study, participants were presented with the scoping framework and an example application of its use using real-world data. Using a pre-defined question set, participants were asked their opinions of key areas of the frameworks and asked to identify any potential limitations with its implementation in practice. The results were, overall, positive with constructive comments for expanding and improving the scoping framework in future work.

This chapter concludes the thesis by first summarising the work conducted, reflecting on the research questions posed in Chapter 1, then discussing how the limitations of the framework, discussed in Chapter 7, can be addressed, alongside an outline of other directions for future work.

8.1 Summary of Research

The motivation behind the work presented in this thesis stems from the challenges that arise for ICS/OT cyber security due to the convergence of IT and OT, otherwise known as Industry 4.0 (Chapter 1). Because of this convergence and increased connectivity in industrial

environments through standardised technologies, a drastic increase in cyber attacks targeting these environments, including CNI, has been observed (Chapter 2).

To this end, an analysis of standards and guidelines and a qualitative study with stakeholders were conducted to assess the effectiveness of current cyber incident R&R in practice (Chapter 3). To address the gaps identified in these studies, a framework is proposed to aid stakeholders in appropriately using standards and guidelines for assessing and improving their cyber incident R&R capabilities (Chapter 4).

Risk management was identified as a crucial phase of the cyber incident R&R lifecycle where adversary-centric security testing can be used as an assurance technique for assessing and improving cyber incident R&R capabilities. However, due to the critical nature of ICS/OT environments, several challenges exist when conducting such engagements in these environments, mainly due to the safety and operational risks present (Chapter 5).

To address this gap, a framework is proposed to aid stakeholders in scoping adversary-centric security tests within ICS/OT environments through quantifying safety and operational risk (Chapter 6). Furthermore, this framework has been evaluated through a qualitative study involving experts in ICS/OT and adversary-centric security testing resulting in positive results for implementing the framework in practice (Chapter 7).

8.2 Reflection on Research Questions

Chapter 1 introduced the key concepts discussed throughout this thesis. ICS/OT is a term to categorise technology involved in controlling or monitoring a physical process, such as a manufacturing or water treatment plants. The importance of securing these systems cannot be understated due to their widespread use within most CNI sectors, such as energy, water and transport. While the convergence of IT and OT through Industry 4.0 has led to several benefits, such as increased process optimisation or remote maintenance and monitoring, this has also led to an increased attack surface for threat actors to target. As such, effectively preparing to defend and respond against cyber attacks targeting ICS/OT has become essential in recent years. Such attacks can severely impact business continuity and safety, such as the 2010 Stuxnet attack on the Iranian Nuclear Enrichment Program [199].

As a means to investigate the described problem space, the following research questions were posited with the goal of identifying current gaps in cyber incident R&R for ICS/OT and devising solutions for addressing these gaps:

- **RQ1**: How has the convergence of IT and OT affected the way that preparation for ICS/OT cyber incident response and recovery needs to be handled?

- **RQ2**: How effective are current ICS/OT cyber incident response and recovery capabilities in practice?
- **RQ3**: Which areas of ICS/OT cyber incident R&R are significantly lacking?
  - **RQ3.1**: How can these areas be improved to better prepare for cyber attacks targeting ICS/OT?
  - **RQ3.2**: Could an approach be developed in these areas to improve cyber incident R&R?

The following sections describe how each research question was addressed individually throughout the research presented in this thesis.

8.2.1 Research Question 1

How has the convergence of IT and OT affected the way that preparation for ICS/OT cyber incident R&R needs to be handled?

Chapter 2 answered this research question in two parts: firstly, through an analysis of cyber attacks targeting ICS/OT and their trends, and, secondly, through an analysis of the differences between IT and OT and how this affects securing industrial environments against cyber threats.

During the analysis of cyber attacks targeting ICS/OT, a total of 43 attacks that occurred between 1988 and 2020 were identified. The evolution of the trends surrounding these attacks demonstrated how Industry 4.0 has led to a shift in how these attacks occur. Before 2009, attacks on ICS/OT were typically carried out by individuals targeting infrastructures where they were employed and motivated by personal reasons. These attacks were mainly disruptive in nature, and initial access was gained through existing levels of access. After 2009, there was a shift towards more organised groups like nation-states and cybercriminals conducting attacks with political motivations like espionage and sabotage, targeting more critical infrastructures like the energy sector. These attacks also relied more on exploiting human vulnerabilities through social engineering for initial access. This shift in trend coincides with the public exposure of the Stuxnet attack in 2010, highlighting the potential damage that attacks on ICS/OT could cause. The convergence of this increased interest in ICS/OT environments from attackers and the shift towards a more interconnected ICS/OT environment contributed to the change in trends observed after 2009.

This rise in attacks and the shift in their trends can be explained through the changes brought about by Industry 4.0, such as increased interconnectivity or the convergence of

Information and Operational Technology. To provide further context on how these changes affect the cyber security of modern ICS/OT environments, an analysis was conducted to compare the fundamental differences between IT and OT. Using asset management standards and guidance, four categories were identified to differentiate IT and OT for analysis: hardware, software, network architecture and protocols, and socio-technical differences. IT hardware is designed to efficiently store, process, and exchange information. In contrast, OT hardware is designed for environmental resilience, high uptime, and cost-efficiency in the monitoring and controlling of operational processes. IT software is flexible and easy to use, while OT software is designed specifically for automation engineers. OT networks prioritise safety over security, leading to the use of protocols that lack basic security features like access control and encryption. These technical differences also contribute to the sociological differences between IT and OT environments, including the implementation of security controls and culture.

8.2.2 Research Question 2

How effective are current ICS/OT cyber incident R&R capabilities in practice?

Two studies were conducted in Chapter 3 to answer this research question. Firstly, an analysis of thirty-one standards and guidelines that provided guidance on cyber incident R&R for ICS/OT environment was conducted. Overall, these resources often reference each other, but there is a lack of consistency in the content and depth of information provided. Technical and procedural processes are both important in cyber incident R&R. However, the analysis found that the reviewed standards and guidelines often lack the necessary level of technical detail or are not tailored to their intended audience. The abundance of resources can make it difficult for operators to select the most comprehensive set of resources and implement processes at the appropriate level of technical depth. This is especially true where the isolated selection of a single resource to drive change within an organisation will likely result in a less than complete picture.

To provide further depth to this analysis of standards and guidelines, a qualitative study was conducted with key stakeholders on implementing these standards in practice. During this study, it was found that most participants had only worked in such sectors and had opportunities to advance into technical and managerial roles. While in-house development of personnel is common, it can lead to isolated viewpoints and a lack of external engagement. Some participants described the development of tailored cyber incident R&R approaches, with a preference for using existing standards and guidelines such as those from NIST. The processes for forming and operating a central incident response team were well understood

and were able to support every potential incident, though the level of internal resources for responding to cyber incidents from an OT perspective was unclear. Documentation of actions during an incident was valued for future use and root cause analysis, and a semi-formal, expert input-based approach was used to evaluate risks during R&R decision-making, prioritising safety but also considering environmental impact, forensic data integrity, and reputational damage. There were conflicting views on the value of standards and guidelines, with some participants finding them helpful in maturing existing processes. In contrast, others believed them to be too information-focused and not function-focused enough for OT. Overall, there was a desire to use existing, proven approaches rather than reinventing the wheel and a need for external engagement and collaboration to provide a holistic approach to cyber security for ICS/OT environments.

8.2.3 Research Question 3

Which areas of ICS/OT cyber incident R&R are significantly lacking?

One solution for the challenges identified in the two studies that were used to answer RQ2 was the proposal of a framework in Chapter 4, which consolidated key standards and guidelines for use for specific phases or sub-phases of the cyber incident R&R lifecycle. However, during the stakeholder interviews in Chapter 3, it was found that the success of certain framework sub-phases in the mid-incident phase of the incident R&R lifecycle depends on the output of other phases. Training of incident response teams was identified as crucial for the success of sub-phases such as incident detection, reporting, containment, eradication, and recovery. Revisiting the 31 standards and guidelines that were analysed in Chapter 3 revealed that sub-phases from the planning and preparation phases contribute highly to the success of sub-phases in the mid-incident and post-incident phases and post-incident sub-phases can contribute to the success of planning and preparation sub-phases for future incidents.

In particular, risk management, which involves identifying, evaluating, and prioritising risks and finding solutions to minimise, monitor, and control them, was found to be crucial due to its effect on the implementation of subsequent phases. Adversary-centric security testing, a method of evaluating and improving cyber resilience and incident response capabilities by emulating the actions of malicious actors, is widespread in IT environments but less so in OT environments. Chapter 5, therefore, answers this research question by providing further depth to the challenges that adversary-centric security testing faces within ICS/OT environments. In this chapter, the technical differences between IT and OT environments were analysed with respect to adversary-centric security testing. Challenges were identified

when conducting these tests in OT environments, particularly in the reconnaissance and weaponisation phases. Passive techniques were found to have little impact on OT operational processes. However, they provided less actionable intelligence, while active techniques returned significant information but had a higher probability of causing operational impact. The use of commonly used IT tools may disrupt the operational process in OT environments, and three main factors when employing these were identified that could adversely affect availability: network throughput of active tools, how resource-intensive these tools are for targets, and the use of unexpected techniques such as vulnerability scans and scripts.

8.2.4 Research Question 3.1

How can these areas be improved to better prepare for cyber attacks targeting ICS/OT?

Despite current approaches that limit the use of active tools during adversary-centric security testing in ICS/OT environments, it is possible to employ such tools subject to the resilience of the systems against more aggressive techniques. Chapter 6, therefore, answers this research question by providing methodologies for quantifying the safety and operational risks of conducting adversary-centric security tests within ICS/OT environments so that these can be better scoped. Indeed, because of the critical nature of these environments, to safely conduct adversary-centric security tests, these risks must be well understood and quantified so that the scoping of these tests can minimise the risks while maximising the depth of testing.

The first step required to quantify safety and operational risks when conducting adversary-centric security tests is to identify hazards that could lead to unwanted impact. (C)HAZOP was identified as a suitable methodology for this as it can be used to identify potential hazards and operability problems that could affect the safe and reliable operation of a system or a process. Once these hazards have been identified, they can be decomposed into smaller, contributing risks, known as basic events, using Fault Tree Analysis which is a deductive technique used to identify and analyse the logical combination of events that can lead to larger hazards (identified using (C)HAZOP) occurring. Data for these basic events can be collected, and the probability of specific tools or techniques that cause a loss of safety or a disruption to the operational process can be precisely quantified. To provide further depth for this, a scenario was devised within a testbed environment that uses real-world hardware to simulate a working industrial environment. Using this scenario, data for employing specific tools and techniques relating to the three identified types of basic events were collected and analysed, allowing for quantifying their risk to safety and the operational process. By understanding the risk that these tools and techniques present, precise scoping constraints

could be defined to ensure that conducting an adversary-centric security test within this environment would not cause a loss of safety or disrupt the operational process.

8.2.5 Research Question 3.2

Could an approach be developed in these areas to improve cyber incident R&R?

With an understanding of how risk can be quantified to better scope adversary-centric security tests within ICS/OT environments to improve cyber incident R&R, a framework is proposed in Chapter 6 to answer this research question. As well as implementing the methodologies discussed for answering RQ3.1, the framework proposes a model based on the Purdue Model and the DiD Model for selecting testing zones for which risk can be identified and quantified. Using these two models to select zones to test, the scoping of adversary-centric security tests can be granularised into separate testing levels with varying degrees of risk to safety and the operational process. Once defined, the safety and operational risks of conducting adversary-centric security tests within these zones can be identified and quantified to define scoping constraints used to reduce the identified risks while maximising the depth of testing. The framework offers a flexible approach to conducting safety and operational risk assessment through example methodologies that can be used at the user's discretion. These example methodologies also give stakeholders appropriate guidance on how to use their current risk assessment procedures for scoping adversary-centric security tests through operational and safety risk quantification.

While the scoping framework is supported by commonly adopted methodologies and data collected using a testbed environment, these only provided a proof of concept of how the framework could be adopted and implemented in practice. Therefore, an evaluative study was conducted in Chapter 7 to assess its applicability in practice and identify potential issues that could limit its use. For this evaluative study, semi-structured interviews were conducted with industry participants that are involved in the field of ICS/OT adversary-centric security testing. During these interviews, participants were provided with the framework and an example application of its use, using data collected from implementing the framework on a testbed-dervied scenario. Following this, they were asked questions regarding each phase of the framework and whether they were appropriate and realistic for implementation in practice. The output of these interviews was then analysed using template analysis for which several themes emerged, covering the phases defined within the scoping framework and other relevant topics. Overall, participants found that the scoping framework successfully addressed the gaps identified throughout the research and could indeed be used to improve

the scoping of adversary-centric security tests within ICS/OT environments to strengthen cyber incident R&R capabilities.

8.3 Future Work

The study conducted in Chapter 7 found that the safety and operational risk-based scoping framework for adversary-centric security tests within ICS/OT environments is a valuable addition to the overall scoping methodology of these engagements. Although a few limitations were identified, they do not affect the framework's validity. Addressing these limitations as part of future work will help to further improve the framework's effectiveness and enhance the scoping of adversary-centric security tests for ICS/OT environments. Overall, the framework was identified as a significant step forward in improving the safety and security of industrial environments.

A few participants during the evaluative study noted that the lack of maturity of users could potentially affect the framework's output and lead to risks not being considered in the scoping process. This lack of maturity has been observed previously throughout the ICS/OT industry as described during the stakeholder engagement conducted in Chapter 3, highlighting the need to drive a cultural change in the way that ICS/OT cyber security is performed. The scoping framework presented in Chapter 6 does begin to address this challenge by providing example methodologies for each phase of the framework; however, further work could be conducted to identify why these cultural challenges exist and how they can be addressed. Notably, concerning the framework itself, integrating standards and guidelines for additional guidance would provide further benefits to framework users that lack the expected level of maturity for its operation.

The study conducted in Chapter 7 identified that further research could be conducted regarding the framework's scalability for large and complex environments. However, several ways exist to address this and ensure that the cost of resources does not outweigh the framework's benefits. For example, through repeated use of the framework, data from previous engagements can be used to enhance the subsequent quantification of risks. This can reduce the need to repeat certain phases, such as risk identification or data collection, thereby saving resources. Another approach to reducing the cost of resources required for the operation of the framework is to leverage automated risk analysis and security testing, a field that has seen considerable research [107, 252, 12]. By integrating automated methods into specific framework phases, such as risk identification or data collection, its efficiency can be improved to reduce the required resources for its operation. With these strategies in

place, the framework can be enhanced to be a valuable tool for scoping adversary-centric security tests within more environments, including large and complex ones.

The critical role of Safety Instrumented Systems in ensuring the safety of industrial environments cannot be overstated. Despite being categorized as OT, their significance in bringing industrial environments back to a safe state in case of any safety hazards makes them a crucial component of these environments. Recognizing their importance, industry standards such as IEC 61511 [117], IEC 61513 [114], and IEC 62061 [120] have been established to provide guidance on the implementation of SISs in specific sectors. Adversary-centric security testing on SISs inherently adds a higher level of risk to safety than devices in levels 0 and 3 of the Purdue Model. Therefore, implementing additional measures for these within the scoping framework would provide additional benefits to users who intend to conduct adversary-centric security tests on their SISs, making the overall testing process more comprehensive and effective.

Lastly, the evaluation conducted in Chapter 7 highlighted the potential benefits of conducting a quantitative evaluation of the framework to further validate its effectiveness. This approach would build on the conducted qualitative evaluation and provide additional confidence for its use in real environments. For this, initial steps have already been taken towards a quantitative evaluation through data collection within the Lancaster University ICS testbed used for the qualitative evaluation of the framework. Conducting further evaluations in real industrial environments would provide tangible evidence of the framework's applicability in practice and further reinforce its value to users.

References

[1] IEEE Standard for Electric Power Systems Communications – Distributed Network Protocol (DNP3). *IEEE Std 1815-2010*, pages 1–775, 2010. doi: 10.1109/IEEESTD.2010.5518537.

[2] AboutSSL. History of the Internet – An Invention That Changed the World. https://aboutssl.org/history-of-the-internet/. Last Accessed: 30-06-2021.

[3] Irfan Ahmed, Sebastian Obermeier, Sneha Sudhakaran, and Vassil Roussev. Programmable logic controller forensics. *IEEE Security & Privacy*, 15(6):18–24, 2017.

[4] Allen Bradley. ControlLogix controllers, revision 16. https://literature.rockwellautomation.com/idc/groups/literature/documents/rn/1756-rn016_-en-e.pdf. Last Accessed: 2022-06-16.

[5] American National Standards Institute / International Society of Automation. ANSI/ISA-5.1-2009 - Instrumentation Symbols and Identification. http://integrated.cc/cse/Instrumentation_Symbols_and_Identification.pdf, 2009.

[6] Moody's Analytics. Saudi Arabia - economic indicators. https://www.economy.com/saudi-arabia/indicators, 2020. Last Accessed: 07-12-2020.

[7] Andrew Ginter. The top 20 cyberattacks on industrial control systems. https://bit.ly/36LfuSe, 2018. Last Accessed: 06-12-2020.

[8] Uchenna P Daniel Ani, Jeremy M Watson, Benjamin Green, Barnaby Craggs, and Jason RC Nurse. Design considerations for building credible security testbeds: Perspectives from industrial control system use cases. *Journal of Cyber Security Technology*, pages 1–49, 2020.

[9] ANSSI. Managing Cybersecurity for Industrial Control Systems. Technical report, Agence Nationale de la Sécurité des Systèmes d'Information, 2012.

[10] Daniele Antonioli, Hamid Reza Ghaeini, Sridhar Adepu, Martin Ochoa, and Nils Ole Tippenhauer. Gamifying ICS security training and research: Design, implementation, and results of S3. In *Proceedings of the 2017 Workshop on Cyber-Physical Systems Security and PrivaCy*, pages 93–102, 2017.

[11] Rob Antrobus, Sylvain Frey, Benjamin Green, and Awais Rashid. SimaticScan: Towards A Specialised Vulnerability Scanner for Industrial Control Systems. In *4th International Symposium for ICS SCADA Cyber Security Research*, pages 11–18, 2016. doi: 10.14236/ewic/ICS2016.2.

[12] Andy Applebaum, Doug Miller, Blake Strom, Chris Korban, and Ross Wolf. Intelligent, automated red team emulation. Technical report, The MITRE Corporation, 2016.

[13] Hilary Arksey and Peter T Knight. *Interviewing for social scientists: An introductory resource with examples*. Sage, London, 1999. ISBN 0761958703.

[14] Liviu Arsene. Oil & gas spearphishing campaigns drop agent tesla spyware in advance of historic OPEC+ deal. https://bit.ly/2YyEcSa, 2020. Last Accessed: 01-07-2020.

[15] Haruna Asai, Tomomi Aoyama, Yuitaka Ota, Yoshihiro Hashimoto, and Ichiro Koshijima. Design and operation framework for industrial control system security exercise. In *Security of Cyber-Physical Systems*, pages 25–51. Springer, 2020.

[16] Michael Assant and Robert Lee. The Industrial Control System Cyber Kill Chain. Technical report, SANS Institute, 2015. Last Accessed: 03-21-2022.

[17] Corline Baylon, Roger Brunt, and David Livingstone. Cyber security at civil nuclear facilities: Understanding the risk. https://bit.ly/2yhfGde, 2015. Last Accessed: 02-07-2020.

[18] Liron Benbenishti. SCADA MODBUS Protocol Vulnerabilities. https://bit.ly/3nEeYy6, 2017. Last Accessed: 15-09-2021.

[19] Molly Betts, Joseph Stirland, Funminiyi Olajide, Kevin Jones, and Helge Janicke. Developing a state of the art methodology & toolkit for ICS SCADA forensics. In *The International Conference on Information Security and Cyber Forensics*, 2016.

[20] Clint Bodungen, Bryan Singer, Aaron Shbeeb, Kyle Wilhoit, and Stephen Hilt. *Hacking Exposed Industrial Control Systems: ICS and SCADA Security Secrets & Solutions*. McGraw Hill Professional, 2016.

[21] Sandro Bologna, Alessandro Fasani, and Maurizio Martellini. Cyber security and resilience of industrial control systems and critical infrastructures. In *Cyber Security*, pages 57–72. Springer, 2013.

[22] William Bolton. *Programmable Logic Controllers*. Elsevier Ltd., sixth edition edition, 2015.

[23] Pauline Bowen, Joan Hash, and Mark Wilson. NIST Special Publication 800-100: Information Security Handbook - A Guide for Managers. Technical report, National Institute of Standards and Technology, 2006.

[24] Stuart A. Boyer. *SCADA: Supervisory Control and Data Acquisition*. ISA, third edition edition, 2004.

[25] Bridewell. Cyber Security: What to Expect in 2023. https://www.bridewell.com/insights/white-papers/detail/cyber-security-what-to-expect-in-2023, 2022. Last Accessed: 20-08-2023.

[26] Jonathan Butts and Michael Glover. How industrial control system security training is falling short. In *International Conference on Critical Infrastructure Protection*, pages 135–149. Springer, 2015.

[27] Eric Byres and Mark Fabro. RISI - The Repository of Industrial Security Incidents. https://www.risidata.com/Database, 2014. Last Accessed: 01-11-2021.

[28] Eric Byres and Dan Hoffman. The myths and facts behind cyber security risks for industrial control systems. In *In Proc. of VDE Kongress*, 2004.

[29] Donald Thomas Campbell and Julian C Stanley. *Experimental and quasi-experimental designs for research on teaching*. Ravenio Books, 1963.

[30] David Canter, Jennifer Brown, and Michael Brenner. *The research interview: Uses and approaches*. Academic Press, New York, 1985. ISBN 0121315800.

[31] Carnegie Mellon University. Network Security Protocols. https://bit.ly/3qNnxX7, 2015.

[32] Carnegie Mellon University. The CERT Division. https://www.sei.cmu.edu/about/divisions/cert/index.cfm, 2019.

[33] Emiliano Casalicchio and Gabriele Gualandi. Asimov: A self-protecting control application for the smart factory. *Future Generation Computer Systems*, 115:213–235, 2021.

[34] Sanjay Chhillar. Common ICS Cybersecurity Myth 1: The Air Gap. https://bit.ly/39JCf9N, 2021. Last Accessed: 28-09-2021.

[35] Chuck Squatriglia. Polish teen hacks his city's trams, chaos ensues. https://bit.ly/2GPiNxs, 2008. Last Accessed: 11-11-2020.

[36] Paul Cichonski, Tom Millar, Tim Grance, and Karen Scarfone. NIST Special Publication 800-61: Computer Security Incident Handling Guide. Technical report, National Institute of Standards and Technology, 2012.

[37] CIS. Cybersecurity spotlight – cyber threat actors. https://www.cisecurity.org/spotlight/cybersecurity-spotlight-cyber-threat-actors/. Last Accessed: 27-11-2020.

[38] CIS. CIS Critical Security Controls. Technical report, Center for Internet Security, 2019.

[39] CISA. CISA - resources. https://us-cert.cisa.gov/resources. Last Accessed: 25-11-2020.

[40] Cisco and Rockwell Automation. Ethernet-to-the-Factory 1.2 Design and Implementation Guide. https://www.cisco.com/c/en/us/td/docs/solutions/Verticals/EttF/EttFDIG.pdf, 2008. Last Accessed: 20-08-2023.

[41] CNSC. REGDOC-2.5.2 Design of Reactor Facilities: Nuclear Power Plants. Technical report, Canadian Nuclear Safety Commission, 2014.

[42] Wm. Arthur Conklin. IT vs OT Security: A Time to Consider a Change in CIA to Include Resilience. In *2016 49th Hawaii International Conference on System Sciences (HICSS)*, pages 2642–2647, 2016. doi: {10.1109/HICSS.2016.331}.

[43] Allan Cook, Helge Janicke, Richard Smith, and Leandros Maglaras. The industrial control system cyber defence triage process. *Computers & Security*, 70:467–481, 2017.

[44] James Cook and Alan Tovey. Honda's global factories brought to a standstill by cyber attack. https://bit.ly/37vqBgD, 2020. Last Accessed: 01-07-2020.

[45] Benjamin F Crabtree and William F Miller. *A Template Approach to Text Analysis: Developing and Using Codebooks*. Sage, Thousand Oaks, CA, US, 1992.

[46] Casey Crane. Recent ransomware attacks: Latest ransomware attack news in 2020. https://www.thesslstore.com/blog/recent-ransomware-attacks-latest-ransomware-attack-news/, 2020. Last Accessed: 26-11-2020.

[47] Jason Creasy and Ian Glover. Cyber Security Incident Response Guide. Technical report, Council for Registered Ethical Security Testers, 2013.

[48] CSO. Petya attack caused $140m hit on Cadbury parent Mondelēz's Q2 revenues. https://bit.ly/3tSMfpn, 2017. Last Accessed: 23-03-2020.

[49] Cybersecurity and Infrastructure Security Agency. Cyber threat source descriptions. https://bit.ly/3rVEsWZ, 2005.

[50] Cybersecurity and Infrastructure Security Agency. Insider threat - cyber. https://www.cisa.gov/insider-threat-cyber, n.d. Last Accessed: 24-05-2021.

[51] E.J. Daniel, C.M. White, and K.A. Teague. An interarrival delay jitter model using multistructure network delay characteristics for packet networks. In *The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1738–1742 Vol.2, 2003. doi: 10.1109/ACSSC.2003.1292282.

[52] Dorothy E Denning. Cyberterrorism: The logic bomb versus the truck bomb. *Global Dialogue*, 2(4):29, 2000.

[53] Dorothy E Denning. Stuxnet: What has changed? *Future Internet*, 4(3):672–687, 2012.

[54] Department of Homeland Security. Emergency Directive 21-01. https://cyber.dhs.gov/ed/21-01/, 2020. Last Accessed: 11-10-2021.

[55] Richard Derbyshire. On the state of OT cyber attacks and traversing level 3.5, the artist formerly known as airgap. https://bit.ly/3JQzbdp, 2023. Last Accessed: 27-03-2023.

[56] Richard Derbyshire, Benjamin Green, Daniel Prince, Andreas Mauthe, and David Hutchison. An analysis of cyber security attack taxonomies. In *2018 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 153–161. IEEE, 2018.

[57] Richard Derbyshire, Benjamin Green, Daniel Prince, Andreas Mauthe, and David Hutchison. An analysis of cyber security attack taxonomies. In *2018 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 153–161. IEEE, 2018.

[58] Richard Derbyshire, Benjamin Green, and David Hutchison. "talking a different language": Anticipating adversary attack cost for cyber risk assessment. *Computers & Security*, 103:102163, 2021. ISSN 0167-4048. doi: https://doi.org/10.1016/j.cose.2020.102163. URL https://www.sciencedirect.com/science/article/pii/S0167404820304363.

[59] Paul Didier, Fernando Macias, James Harstad, Rick Antholine, Scott A. Johnston, Sabina Piyevsky, Dan Zaniewski, Steve Zuponcic, Mark Schillace, and Gregory Wilcox. Converged Plantwide Ethernet (CPwE) Design and Implementation Guide. *Rockwell Automation*, 9:564, 2011.

[60] Marietheres Dietz, Manfred Vielberth, and Günther Pernul. Integrating digital twin security simulations in the security operations center. In *Proceedings of the 15th International Conference on Availability, Reliability and Security*, ARES '20, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450388337. doi: 10.1145/3407023.3407039. URL https://doi.org/10.1145/3407023.3407039.

[61] Ben Dooley and Hisako Ueon. Honda hackers may have used tools favored by countries. https://nyti.ms/2BdmRUU, 2020. Last Accessed: 01-07-2020.

[62] Dragos. Crashoverride: Analysis of the threat to electric grid operations. https://bit.ly/2KxWN8r, 2017. Last Accessed: 02-07-2020.

[63] Dragos. ICS/OT Cybersecurity Year in Review. https://hub.dragos.com/ics-cybersecurity-year-in-review-2022, 2022. Last Accessed: 20-08-2023.

[64] David P. Duggan. Penetration Testing of Industrial Control Systems. https://bit.ly/3AkNeCm, 2005. Last Accessed: 16-09-2021.

[65] Jordi Dunjó, Vasilis Fthenakis, Juan A. Vílchez, and Josep Arnaldos. Hazard and Pperability (HAZOP) Analysis. A Literature Review. *Journal of Hazardous Materials*, 173(1):19–32, 2010. ISSN 0304-3894. doi: https://doi.org/10.1016/j.jhazmat.2009.08.076. URL https://www.sciencedirect.com/science/article/pii/S0304389409013727.

[66] DWI. Guidance on the Implementation of the NIS Regulations 2018 - The Cyber Assessment Framework (CAF). Technical report, Drinking Water Inspectorate, 2018.

[67] DWI. CAF Information. http://dwi.defra.gov.uk/nis/caf/index.html, 2019.

[68] EC-Council. Certified Ethical Hacker Certification. https://bit.ly/3cK7t23, 2021. Last Accessed: 24-11-2021.

[69] Peter Eden, Andrew Blyth, Pete Burnap, Yulia Cherdantseva, Kevin Jones, Hugh Soulsby, and Kristan Stoddart. Forensic readiness for SCADA/ICS incident response. In *4th International Symposium for ICS & SCADA Cyber Security Research 2016 4*, pages 142–150, 2016.

[70] Dmitry Efanov. PLCScan. https://code.google.com/archive/p/plcscan/, 2012. Last Accessed: 2022-06-27.

[71] ENISA. Good Practice Guide for Incident Management. Technical report, European Network and Information Security Agency, 2010.

[72] ENISA. The NIS Directive. https://bit.ly/3D0Bo27, 2018. Last Accessed: 15-06-2021.

[73] Sead Fadilpašić. Agent tesla malware receives module for stealing wi-fi passwords. https://bit.ly/3b1qwl2, 2020. Last Accessed: 01-07-2020.

[74] James P Farwell and Rafal Rohozinski. Stuxnet and the future of cyber war. *Survival*, 53(1):23–40, 2011.

[75] Davide Fauri, Bart de Wijs, Jerry den Hartog, Elisa Costante, Emmanuele Zambon, and Sandro Etalle. Encryption in ICS networks: A blessing or a curse? In *2017 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 289–294, 2017. doi: 10.1109/SmartGridComm.2017.8340732.

[76] Federal Bureau of Investigation. Welcome to infragard. https://www.infragard.org/, 2020. Last Accessed: 17-12-2020.

[77] Barbara Filkins and Doug Wylie. SANS 2019 State of OT/ICS Cybersecurity Survey. https://bit.ly/3rP9amY, 2019.

[78] Fortinet. 2022 State of Operational Technology and Cybersecurity Report. https://www.fortinet.com/content/dam/fortinet/assets/analyst-reports/report-2022-ot-cybersecurity.pdf, 2022. Last Accessed: 20-08-2023.

[79] Fortinet. The CIA Triad. https://bit.ly/3CDsEOK, n.d. Last Accessed: 01-11-21.

[80] Ivo Friedberg, Kieran McLaughlin, Paul Smith, David Laverty, and Sakir Sezer. STPA-SafeSec: Safety and security analysis for cyber-physical systems. *Journal of Information Security and Applications*, 34:183–196, 2017. ISSN 2214-2126. doi: https://doi.org/10.1016/j.jisa.2016.05.008. URL https://www.sciencedirect.com/science/article/pii/S2214212616300850.

[81] Heinz Gall. Functional safety IEC 61508 / IEC 61511 the impact to certification and the user. In *2008 IEEE/ACS International Conference on Computer Systems and Applications*, pages 1027–1031, 2008. doi: 10.1109/AICCSA.2008.4493673.

[82] Deianira Ganga and Sam Scott. Cultural "insiders" and the issue of positionality in qualitative migration research: Moving "across" and moving "along" researcher-participant divides. In *Forum Qualitative Sozialforschung/Forum: Qualitative Social Research*, volume 7, 2006. ISBN 1438-5627.

[83] Joseph Gardiner, Barnaby Craggs, Benjamin Green, and Awais Rashid. Oops i did it again: Further adventures in the land of ICS security testbeds. CPS-SPC'19, page 75–86, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450368315. doi: 10.1145/3338499.3357355. URL https://doi.org/10.1145/3338499.3357355.

[84] Gartner. Gartner glossary - operational technology (OT). https://www.gartner.com/en/information-technology/glossary/operational-technology-ot, n.d. Last Accessed: 08-11-2022.

[85] Claire Gaving. Seasonal Variations in Electricity Demand. https://bit.ly/3qustQv, 2014.

[86] GIAC. GIAC Penetration Tester (GPEN). https://bit.ly/3G00Rt9, 2021. Last Accessed: 24-11-2021.

[87] Martin Giles. Triton is the world's most murderous malware, and it's spreading. https://bit.ly/2W0Bmm0, 2019. Last Accessed: 01-07-2020.

[88] Barney Glaser and Anselm Strauss. Grounded theory: the discovery of grounded theory. *Sociology The Journal Of The British Sociological Association*, 12:27–49, 1967.

[89] Fernando Gont. Security Assessment of the Internet Protocol. https://bit.ly/3dt2gwr, 2008. Last Accessed: 20-08-2023.

[90] Jesus Gonzalez and Mauricio Papa. Passive scanning in modbus networks. In *International Conference on Critical Infrastructure Protection*, pages 175–187. Springer, 2007.

[91] Benjamin Green, Marina Krotofil, and Ali Abbasi. On the significance of process comprehension for conducting targeted ICS attacks. In *Proceedings of the 2017 Workshop on Cyber-Physical Systems Security and PrivaCy*, pages 57–67. ACM, 2017.

[92] Benjamin Green, Anhtuan Lee, Rob Antrobus, Utz Roedig, David Hutchison, and Awais Rashid. Pains, Gains and PLCs: Ten Lessons from Building an Industrial Control Systems Testbed for Security Research. In *10th USENIX Workshop on Cyber Security Experimentation and Test (CSET 17)*, 2017.

[93] Benjamin Green, Daniel Prince, Jerry Busby, and David Hutchison. "How Long is a Piece of String": Defining Key Phases And Observed Challenges within ICS Risk Assessment. New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450353946. doi: 10.1145/3140241.3140251. URL https://doi.org/10.1145/3140241.3140251.

[94] Benjamin Green, Richard Derbyshire, William Knowles, James Boorman, Pierre Ciholas, Daniel Prince, and David Hutchison. ICS Testbed Tetris: Practical Building Blocks Towards a Cyber Security Resource. In *13th USENIX Workshop on Cyber Security Experimentation and Test (CSET 20)*, 2020.

[95] Benjamin Green, Richard Derbyshire, Marina Krotofil, William Knowles, Daniel Prince, and Neeraj Suri. PCaaD: Towards Automated Determination and Exploitation of Industrial Systems. *Computers & Security*, 110:102424, 2021.

[96] Andy Greenberg. *Sandworm: A New Era of Cyberwar and the Hunt for the Kremlin's Most Dangerous Hackers*. Doubleday, 2019.

[97] guru99. Real-time operating system (RTOS): Components, Types, Examples. https://www.guru99.com/real-time-operating-system.html. Last Accessed: 28-06-2021.

[98] Amin Hassanzadeh, Amin Rasekh, Stefano Galelli, Mohsen Aghashahi, Riccardo Taormina, Avi Ostfeld, and M Katherine Banks. A review of cybersecurity incidents in the water sector. *Journal of Environmental Engineering*, 146(5):03120003, 2020.

[99] Ying He, Leandros A Maglaras, Helge Janicke, and Kevin Jones. An industrial control systems incident response decision framework. In *2015 IEEE Conference on Communications and Network Security (CNS)*, pages 761–762. IEEE, 2015.

[100] Kevin E Hemsley, E Fisher, et al. History of industrial control system cyber incidents. Technical report, Idaho National Lab.(INL), Idaho Falls, ID (United States), 2018.

[101] Mariana Hentea. *Building an Effective Security Program for Distributed Energy Resources and Systems*. John Wiley & Sons Inc., 2021.

[102] Gregg Herken. Thomas c. reed, at the abyss: An insider's history of the cold war. new york: Ballantine books, 2004. 368 pp, 2007.

[103] Hidekazu Hirai, TOMOMI Aoyama, N Davaadorj, and ICHIRO Koshijima. Framework for cyber incident response training. *Safety and Security Engineering VII, Rome, Italy*, pages 273–283, 2017.

[104] HMG. Security Policy Framework. https://www.gov.uk/government/publications/security-policy-framework/hmg-security-policy-framework, 2022. Last Accessed 20-08-2023.

[105] Ivan Homoliak, Flavio Toffalini, Juan Guarnizo, Yuval Elovici, and Martín Ochoa. Insight into insiders and it: A survey of insider threat taxonomies, analysis, modeling, and countermeasures. Technical report, Singapore University of Technology and Design, 2019.

[106] HSE. Cyber Security for Industrial Automation and Control Systems (IACS). http://www.hse.gov.uk/foi/internalops/og/og-0086.pdf, 2018.

[107] Tanvir Hussain and Robert Eschbach. Automated fault tree generation and risk-based testing of networked automation systems. In *2010 IEEE 15th Conference on Emerging Technologies & Factory Automation (ETFA 2010)*, pages 1–8, 2010. doi: 10.1109/ETFA.2010.5641309.

[108] Eric Hutchins, Michael Cloppert, and Rohan Amin. Intelligence-Driven Computer Network Defense and Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains. Technical report, Lockheed Martin. URL https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/LM-White-Paper-Intel-Driven-Defense.pdf.

[109] IAEA. IAEA Nuclear Security Series No. 17. Technical report, International Atomic Energy Agency, 2011.

[110] IAEA. Nuclear Security Fundamentals: Objective and Essential Elements of a State's Nuclear Security Regime. Technical report, International Atomic Energy Agency, 2013.

[111] IAEA. IAEA Nuclear Security Series No. 23-G. Technical report, International Atomic Energy Agency, 2015.

[112] ICS-CERT. ICS Advisory (ICSA-14-178-01): ICS Focused Malware. https://bit.ly/356F5CM, 2014. Last Accessed: 02-07-2020.

[113] IEC. IEC 61025:2006 - Fault Tree Analysis, 2006.

[114] IEC. IEC 61513:2011: Nuclear Power Plants - Instrumentation and Control Important to Safety, 2011.

[115] IEC. BS EN IEC 62443-2-1:2011, 2011.

[116] IEC. IEC 61131-3:2013, 2013.

[117] IEC. IEC 61511-1:2016: Functional Safety - Safety Instrumented Systems for the Process Industry Sector, 2016.

[118] IEC. IEC 62443, 2019.

[119] IEC. BS EN IEC 62443-4-2:2019, 2019.

[120] IEC. IEC 62061:2021: Safety of Machinery - Functional Safety of Safety-Related Control Systems, 2021.

[121] PTC Inc. ThingWorx IIoT Platform, 2019.

[122] International Atomic Energy Agency. Nuclear Share of Electricity Generation in 2019. https://pris.iaea.org/ pris/worldstatistics/ nuclearshareofelectricitygeneration.aspx, 2021.

[123] International Society of Automation. ISA Courses. https://www.isa.org/store, 2021. Last Accessed: 24-11-2021.

[124] International Society of Automation. Overview of Penetration Testing for Industrial Control Systems (IC38C). https://bit.ly/2ZhvhqO, 2021. Last Accessed: 24-11-2021.

[125] ISO. ISO 14617-6:2002 - Graphical symbols for diagrams Part 6: Measurement and control functions, 2002.

[126] ISO. ISO 31000:2018 - Risk management — Guidelines, 2018.

[127] ISO/IEC. BS ISO/IEC 27035-1:2016, 2016.

[128] ISO/IEC. BS ISO/IEC 27035-2:2016, 2016.

[129] ISO/IEC. BS EN ISO/IEC 27001:2017, 2017.

[130] ISO/IEC. BS EN ISO/IEC 27002:2017, 2017.

[131] ISO/IEC. BS EN ISO/IEC 27019:2017, 2017.

[132] Martin Gilje Jaatun, Eirik Albrechtsen, Maria B Line, Inger Anne Tøndel, and Odd Helge Longva. A framework for incident response management in the petroleum industry. *International Journal of Critical Infrastructure Protection*, 2(1-2):26–37, 2009.

[133] William Jardine, Sylvain Frey, Benjamin Green, and Awais Rashid. Senami: Selective non-invasive active monitoring for ICS intrusion detection. In *Proceedings of the 2nd ACM Workshop on Cyber-Physical Systems Security and Privacy*, pages 23–34, 2016.

[134] Joby Warrick, Ellen Nakashima. Official: Israel linked to a disruptive cyber-attack on iranian port facility. https://wapo.st/3kpjGKS, 2020. Last Accessed: 11-11-2020.

[135] John Leyden. Polish teen derails tram after hacking train network. https://bit.ly/3eKROzK, 2008. Last Accessed: 11-11-2020.

[136] Blake Johnson, Dan Caban, Marina Krotofil, Dan Scali, Nathan Brubaker, and Christopher Glyer. Attackers deploy new ICS attack framework "TRITON" and cause operational disruption to critical infrastructure. *Threat Research Blog*, 14, 2017.

[137] Kevin Jornson and Mike Snider. Pompeo Says Russia "Pretty Clearly" Behind Cyber-attack on US, but Trump Casts Doubts and Downplays Threat. https://bit.ly/2YGFg8k, 2020. Last Accessed: 11-10-2021.

[138] Mehmet Karakas. Determination of Network Delay Distribution over the Internet. https://etd.lib.metu.edu.tr/upload/1223155/index.pdf, 2003. Last Accessed: 20-08-2023.

[139] Frank Kargl, Rens W van der Heijden, Hartmut König, Alfonso Valdes, and Marc C Dacier. Insights on the security and dependability of industrial control systems. *IEEE security & privacy*, 12(6):75–78, 2014.

[140] Kaspersky ICS CERT. H1 2022 - a brief overview of the main incidents in industrial cybersecurity. https://bit.ly/3FUI0Sv, 2022. Last Accessed: 20-08-2023.

[141] Anastasis Keliris, Hossein Salehghaffari, Brian Cairl, Prashanth Krishnamurthy, Michail Maniatakos, and Farshad Khorrami. Machine learning-based defense against process aware attacks on industrial control systems. Technical report, New York University, 2016.

[142] David Kennedy. The Social Engineer Toolkit. https://bit.ly/3CsnGDR, 2020. Last Accessed: 20-09-2021.

[143] Abdullah Khalili, Ashkan Sami, Mahsa Keikha, and Ali Akbar Safavi. Recovery scheme for industrial control systems. In *The 5th Conference on Information and Knowledge Technology*, pages 279–283. IEEE, 2013.

[144] Joonsoo Kim, Kyeongho Kim, and Moonsu Jang. Cyber-physical battlefield platform for large-scale cybersecurity exercises. In *2019 11th International Conference on Cyber Conflict (CyCon)*, volume 900, pages 1–19. IEEE, 2019.

[145] Nigel King, C Cassell, and G Symon. Qualitative methods in organizational research: A practical guide. *The Qualitative Research Interview*, 17, 1994.

[146] T.A. Kletz and Institution of Chemical Engineers (Great Britain). *HAZOP & HAZAN: Notes on the Identification and Assessment of Hazards*. Hazard workshop modules. Institution of Chemical Engineers, 1986. ISBN 9780852951651. URL https://books.google.fr/books?id=RrGUQgAACAAJ.

[147] Eric D. Knapp and Joel Thomas Langill. *Industrial Network Security*. Elsevier Ltd., 2015.

[148] William Knowles, Jose M. Such, Antonios Gouglidis, Gaurav Misra, and Awais Rashid. Assurance Techniques for Industrial Control Systems (ICS). New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450338271. doi: 10.1145/2808705.2808710. URL https://doi.org/10.1145/2808705.2808710.

[149] Eduard Kovacs. Israel says hackers targeted SCADA systems at water facilities. https://bit.ly/3fiHDlQ, 2020. Last Accessed: 01-07-2020.

[150] Eduard Kovacs. Hackers knew how to target PLCs in israel water facility attacks: Sources. https://bit.ly/35vwNnQ, 2020. Last Accessed: 01-07-2020.

[151] Patrick Kral. Information Security Reading Room: Incident Handler's Handbook. https://www.sans.org/reading-room/whitepapers/incident/incident-handlers-handbook-33901, 2019. Last Accessed: 20-08-2023.

[152] Lancaster University. ICS Response and Recovery Framework. https://ics-rr-framework.github.io/, 2021.

[153] Ralph Langner. Stuxnet: Dissecting a Cyberwarfare Weapon. *IEEE Security & Privacy*, 9(3):49–51, 2011.

[154] Robert M Lee, Michael J Assante, and Tim Conway. German Steel Mill Cyber Attack. *Industrial Control Systems*, 30:62, 2014.

[155] Nancy G. Leveson. *Engineering a Safer World: Systems Thinking Applied to Safety*. The MIT Press, 01 2012. ISBN 9780262298247. doi: 10.7551/mitpress/8179.001.0001. URL https://doi.org/10.7551/mitpress/8179.001.0001.

[156] John Leyden. Uk teenager accused of 'electronic sabotage' against us port. https://bit.ly/2ymdVes, 2003. Last Accessed: 02-07-2020.

[157] John Leyden. Mystery lingers over stealthy stuxnet infection. https://bit.ly/2RQ5mzD, 2010. Last Accessed: 01-07-2020.

[158] Maria B Line, Ali Zand, Gianluca Stringhini, and Richard Kemmerer. Targeted attacks against industrial control systems: Is the power industry prepared? In *Proceedings of the 2nd Workshop on Smart Energy Grid Security*, pages 13–22, 2014.

[159] Maria Bartnes Line, Inger Anne Tøndel, and Martin G Jaatun. Current practices and challenges in industrial control organizations regarding information security incident management–does size matter? information security incident management in large and small industrial control organizations. *International journal of critical infrastructure protection*, 12:12–26, 2016.

[160] Jonathan Love. *Hazard Analysis - Process Automation Handbook: A Guide to Theory and Practice*. Springer London, 2007. ISBN 978-1-84628-282-9. doi: 10.1007/978-1-84628-282-9_54. URL https://doi.org/10.1007/978-1-84628-282-9_54.

[161] Gordon Lyon. Nmap: the Network Mapper. https://nmap.org/, 1997. Last Accessed: 15-09-2021.

[162] Sam Maesschalck, Alexander Staves, Richard Derbyshire, Benjamin Green, and David Hutchison. Walking under the ladder logic: PLC-VBS: a PLC control logic vulnerability scanning tool. *Computers & Security*, 127:103116, 2023. ISSN 0167-4048. doi: https://doi.org/10.1016/j.cose.2023.103116. URL https://www.sciencedirect.com/science/article/pii/S0167404823000263.

[163] Kevin Maguire. Guard tried to sabotage nuclear reactor. https://bit.ly/3cd4f4S, 2001. Last Accessed: 02-07-2020.

[164] Malwarebytes. Honda and Enel impacted by cyber attack suspected to be ransomware. https://bit.ly/2YF9Dsa, 2020. Last Accessed: 01-07-2020.

[165] John Matherly. Shodan Search Engine. https://www.shodan.io/, 2009. Last Accessed: 15-09-2021.

[166] McAfee. Global energy cyberattacks: Night dragon. https://bit.ly/2RN5JL8, 2011. Last Accessed: 22-06-2020.

[167] Microsoft. Windows Comprehensive Security. https://bit.ly/3qnouGx, . Last Accessed: 26-06-2021.

[168] Microsoft. Windows Server. https://bit.ly/3CXOspc, . Last Accessed: 26-06-2021.

[169] Microsoft. Microsoft security intelligence: Virus:win32/sality.at. https://bit.ly/3ahhznu, 2010. Last Accessed: 22-06-2020.

[170] Bill Miller and Dale C Rowe. A survey scada of and critical infrastructure incidents. *RIIT*, 12:51–56, 2012.

[171] Thomas Miller, Alexander Staves, Sam Maesschalck, Miriam Sturdee, and Benjamin Green. Looking Back to Look Forward: Lessons learnt from Cyber-Attacks on Industrial Control Systems. *International Journal of Critical Infrastructure Protection*, 35:100464, 2021. ISSN 1874-5482. doi: https://doi.org/10.1016/j.ijcip.2021.100464. URL https://www.sciencedirect.com/science/article/pii/S1874548221000524.

[172] Gyorgy Miru. Siemens S7 Communication - Part 1 General Structure. http://gmiru.com/article/s7comm/, 2017. Last Accessed: 08-07-2021.

[173] MITRE. ATT&CK ICS - Overview. https://collaborate.mitre.org/attackics/index.php/Overview, 2020. Last Accessed: 07-12-2020.

[174] MITRE. ATT&CK Matrix for Enterprise. https://bit.ly/3ggl94C, 2020. Last Accessed: 03-07-2020.

[175] MITRE. MITRE Common Vulnerability and Exposures. https://cve.mitre.org/, 2021. Last Accessed: 21-09-2021.

[176] MITRE. ATT&CK Matrix for Enterprise. https://attack.mitre.org/, 2021. Last Accessed: 22-12-2021.

[177] MITRE. Adversary emulation plans, 2022. URL https://attack.mitre.org/resources/adversary-emulation-plans/. Last Accessed: 24-09-2022.

[178] MITRE. ATT&CK for Industrial Control Systems. https://attack.mitre.org/matrices/ics/, 2022.

[179] Mitsubishi Electric. MELSEC Communication Protocol Reference Manual. https://dl.mitsubishielectric.com/dl/fa/document/manual/plc/sh080008/sh080008ab.pdf, n.d. Last Accessed: 20-08-2023.

[180] Elizabeth Montalbano. U.k. water supplier hit with clop ransomware attack. https://threatpost.com/water-supplier-hit-clop-ransomware/180422/, 2022. Last Accessed: 27-03-2023.

[181] G.E. Moore. Cramming more components onto integrated circuits. volume 86, pages 82–85, 1998. doi: 10.1109/JPROC.1998.658762.

[182] Dan Morain. Hackers victimize Cal-ISO. https://lat.ms/3bfitC05, 2001. Last Accessed: 02-07-2020.

[183] Mykhailo Mozhaiev, Nina Kuchuk, and Maksym Usatenko. The method of jitter determining in the telecommunication network of a computer system on a special software platform. *Innovative Technologies and Scientific Solutions for Industries*, pages 134–140, 12 2019. doi: 10.30837/2522-9818.2019.10.134.

[184] Glenn Murray, Michael N. Johnstone, and Craig Valli. The Convergence of IT and OT in Critical Infrastructure. 2017.

[185] National Cyber Security Centre. NCSC CAF Guidance. https://www.ncsc.gov.uk/collection/caf, 2021.

[186] National Cyber Security Centre. About the NCSC. https://www.ncsc.gov.uk/section/about-ncsc/what-we-do, 2021.

[187] National Institute of Standards and Technology. Framework for Improving Critical Infrastructure Cybersecurity. https://bit.ly/369NPgo, 2018. Last Accessed: 20-08-2023.

[188] National Institute of Standards and Technology. NIST Special Publication 1800-5. https://bit.ly/3w3MFd7, 2018. Last Accessed: 20-08-2023.

[189] National Security Agency. Defense in Depth: A practical strategy for achieving Information Assurance in today's highly networked environments. https://web.archive.org/web/20121002051613/https://www.nsa.gov/ia/_files/support/defenseindepth.pdf, 2012. Last Accessed: 12/07/2022.

[190] NCSC. About the ncsc - what we do. https://www.ncsc.gov.uk/section/about-ncsc/what-we-do, . Last Accessed: 25-11-2020.

[191] NCSC. NCSC - search results: Phishing. https://www.ncsc.gov.uk/search?q=phishing&start=0&rows=20&articleType=guidance, . Last Accessed: 26-11-2020.

[192] NCSC. 10 Steps to Cyber Security: Incident Management. https://www.ncsc.gov.uk/collection/10-steps-to-cyber-security?curPage=/collection/10-steps-to-cyber-security/the-10-steps/incident-management, 2018.

[193] NCSC. CHECK - Penetration Testing. https://bit.ly/3xfOIx0, 2019.

[194] NCSC. NCSC CAF Guidance. https://www.ncsc.gov.uk/collection/nis-directive?curPage=/collection/nis-directive/introduction-to-the-nis-directive, 2019.

[195] NCSC. NCSC CAF Guidance Objective C - Detecting Cyber Security Events. https://www.ncsc.gov.uk/collection/caf/cyber-assessment-framework/caf-objective-c-detecting-cyber-security-events, 2019.

[196] Mar Negreiro. The NIS2 Directive - A high common level of cybersecurity in the EU. https://www.europarl.europa.eu/RegData/etudes/BRIE/2021/689333/EPRS_BRI(2021)689333_EN.pdf, 2022. Last Accessed: 20-08-2023.

[197] NERC. CIP-008-6 - Cyber Security - Incident Reporting and Response Planning. Technical report, North American Electric Reliability Corporation, 2019.

[198] Andrew Nicholson, Stuart Webber, Shaun Dyer, Tanuja Patel, and Helge Janicke. SCADA security in the light of cyber-warfare. *Computers & Security*, 31(4):418–436, 2012.

[199] Falliere Nicolas, Liam Murchu, and Eric Chien. W32.Stuxnet Dossier Version 1.3. https://bit.ly/3aqefqg, 2010. Last Accessed: 14-06-2021.

[200] Jakob Nielsen and Thomas K. Landauer. A mathematical model of the finding of usability problems. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems*, page 206–213. Association for Computing Machinery, 1993. URL https://doi.org/10.1145/169059.169166.

[201] NIST. NIST special publication 800-series general information. https://www.nist.gov/itl/publications-0/nist-special-publication-800-series-general-information. Last Accessed: 02/12/2020.

[202] NIST. Draft NIST Special Publication 800-53, Revision 5, Initial Public Draft. https://csrc.nist.gov/csrc/media/publications/sp/800-53/rev-5/draft/documents/sp800-53r5-draft.pdf, 2017. Last Accessed: 20-08-2023.

[203] NIST. The National Vulnerability Database. https://nvd.nist.gov/, 2021. Last Accessed: 21-09-2021.

[204] NITTF. Insider threat program - maturity framework. https://bit.ly/2NkcqFD, 2018. Last Accessed: 15-03-2020.

[205] NPSA. Critical National Infrastructure. https://bit.ly/3ueRgu6, 2021. Last Accessed: 20-08-2023.

[206] NRC. RG 5.71 Cyber Security Programs for Nuclear Facilities. https://www.nrc.gov/docs/ML0903/ML090340159.pdf, 2010. Last Accessed: 20-08-2023.

[207] Nuclear Energy Institute. NEI 08-09 Cyber Security Plan for Nuclear Power Reactors. https://www.nrc.gov/docs/ML1011/ML101180437.pdf, 2010. Last Accessed: 20-08-2023.

[208] OASIS. Introduction to STIX. https://oasis-open.github.io/cti-documentation/stix/intro, 2020. Last Accessed: 24-11-2020.

[209] ODVA. ODVA Specifications. https://bit.ly/3isYlSl, n.d. Last Accessed: 08-07-2021.

[210] Offensive Security. Courses and Certifications, 2021. URL https://www.offensive-security.com/courses-and-certifications/. Last Accessed: 14-06-2021.

[211] Office for Budget Responsibility. Cyber-attacks during the russian invasion of ukraine. https://obr.uk/box/cyber-attacks-during-the-russian-invasion-of-ukraine/, 2022. Last Accessed: 27-03-2022.

[212] Office for Nuclear Regulation. Security Assessment Principles for the Civil Nuclear Industry. http://www.onr.org.uk/syaps/security-assessment-principles-2017.pdf, 2017. Last Accessed: 20-08-2023.

[213] Office of the Press Secratary. Executive Order - Improving Critical Infrastructure Cybersecurity. https://obamawhitehouse.archives.gov/the-press-office/2013/02/12/executive-order-improving-critical-infrastructure-cybersecurity, 2021.

[214] ONR. Office for Nuclear Regulation (ONR) Permissioning Inspection - Technical Assessment Guides. http://www.onr.org.uk/operational/tech_asst_guides/, 2019.

[215] Dan Palade, Charles Møller, Chen Li, and Soujanya Mantravadi. An open platform for smart production: IT/OT integration in a smart factory. In *ICEIS (2)*, pages 707–714, 2021.

[216] Michael Quinn Patton. *Qualitative evaluation and research methods*. SAGE, London, 1990. ISBN 0803937792.

[217] PenTestPartners. Introduction to PLCs and Ladder Logic. https://bit.ly/3hoxXIv, . Last Accessed: 26-06-2021.

[218] PenTestPartners. Snakes and Ladder Logic. https://bit.ly/2U3sz5x, . Last Accessed: 26-06-2021.

[219] Peplink. MAX Outdoor Router. https://www.peplink.com/products/max-cellular-router/outdoor/{#}hd2, 2019.

[220] Ricardo Silva Peres, Xiaodong Jia, Jay Lee, Keyi Sun, Armando Walter Colombo, and Jose Barata. Industrial artificial intelligence in industry 4.0 - systematic review, challenges and outlook. *IEEE Access*, 2020. doi: 10.1109/ACCESS.2020.3042874.

[221] Nicole Perlroth. Hackers target operator of kansas nuclear power plant, fbi and homeland security say. https://bit.ly/2yLYht4, 2017. Last Accessed: 01-07-2020.

[222] Andrés F Murillo Piedrahita, Vikram Gaur, Jairo Giraldo, Alvaro A Cardenas, and Sandra Julieta Rueda. Leveraging software-defined networking for incident response in industrial control systems. *IEEE Software*, 35(1):44–50, 2017.

[223] Janet Powney and Mike Watts. *Interviewing in educational research*. Routledge & Kegan Paul, Ablingdon, 1987. ISBN 0710206232.

[224] Proofpoint. Emerging threats intelligence. https://www.proofpoint.com/us/products/advanced-threat-protection/et-intelligence, 2020. Last Accessed: 17-12-2020.

[225] PTC Inc. KEPServerEX - Solving Your Communications Challenges. https://www.kepware.com/en-us/products/, 2019.

[226] Rapid7. The Metasploit Framework. https://bit.ly/3AFSVdW, 2003. Last Accessed: 21-09-2021.

[227] Gelli Ravikumar, Burhan Hyder, and Manimaran Govindarasu. Next-generation CPS testbed-based grid exercise-synthetic grid, attack, and defense modeling. In *2020 Resilience Week (RWS)*, pages 92–98. IEEE, 2020.

[228] Rebecca Addison. Israel linked to cyber-attack on iranian port. https://bit.ly/35lOoQR, 2020. Last Accessed: 11-11-2020.

[229] Weber RI. *Basic Content Analysis*. Beverly Hills, CA: Sage Publications, 1985.

[230] Kai Roer. The security culture framework. https://securityculturеframework.net/, 2021. Last Accessed: 25-05-2021.

[231] Herbert J Rubin and Irene S Rubin. *Qualitative interviewing: The art of hearing data*. Sage, London, 2011. ISBN 1452285861.

[232] Emmanouil Samanis, Joseph Gardiner, and Awais Rashid. SoK: A taxonomy for contrasting industrial control systems asset discovery tools. In *Proceedings of the 17th International Conference on Availability, Reliability and Security*, ARES '22, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450396707. doi: 10.1145/3538969.3538979. URL https://doi.org/10.1145/3538969.3538979.

[233] Alex Samonte. Network Security Reference Architecture. https://bit.ly/3Ahn5o7. Last Accessed: 30-06-2021.

[234] SANS Institute. Cybersecurity Courses & Certifications. https://bit.ly/3nLLZYQ, 2021. Last Accessed: 24-11-2021.

[235] SANS Institute. Assessing and Exploiting Control Systems. https://bit.ly/3oVGhTM, 2021. Last Accessed: 24-11-2021.

[236] Tsubasa Sasaki, Kenji Sawada, Seiichi Shin, and Shu Hosokawa. Fallback and recovery control system of industrial control system for cybersecurity. *IFAC-PapersOnLine*, 50(1):15247–15252, 2017.

[237] Schneider Electric. ClearSCADA: Software for Telemetry and Remote SCADA Systems and Applications. https://www.schneider-electric.co.uk/en/product-range-presentation/61264-clearscada/, 2019.

[238] Schneider Electric. SCADAPack 100, 300, 32. https://bit.ly/3zeIAGY, 2019. Last Accessed: 25-03-2023.

[239] Klaus Schwab. *The fourth industrial revolution*. Currency, 2017.

[240] Tara Seals. SAS 2019: Triton ICS malware hits a second victim. https://bit.ly/3eVvrqP, 2019. Last Accessed: 01-07-2020.

[241] Justine Searle. Control things i/o. https://www.controlthings.io/home, 2021. Last Accessed: 27-06-2022.

[242] Offensive Security. Kali linux. https://www.kali.org/, 2022. Last Accessed: 25-09-2022.

[243] Panda Security. Mariposa botnet. https://bit.ly/3bprp8e, 2010. Last Accessed: 01-07-2020.

[244] Panda Security. The most famous virus in history: Blaster. https://bit.ly/2Vh88jJ, 2014. Last Accessed: 02-07-2020.

[245] Panda Security. Critical infrastructure. https://bit.ly/2VJtPrJ, 2018. Last Accessed: 02-07-2020.

[246] Chris Sherry. Advantages and Disadvantages of Active vs. Passive Scanning in IT and OT Environments, 2020. URL https://bit.ly/3trOgLy. Last Accessed: 14-06-2021.

[247] Siemens. SIMATIC S7-1200. https://new.siemens.com/global/en/products/automation/systems/industrial/plc/s7-1200.html. Last Accessed: 2022-06-16.

[248] Siemens. KTP700F. https://mall.industry.siemens.com/mall/en/uk/Catalog/Product/6AV2125-2GB23-0AX0, 2019.

[249] Siemens. ET200S. https://sie.ag/3TRz0TU, 2019.

[250] Siemens. SIMATIC S7-300 - Proven Multiple Times!, 2020. URL https://sie.ag/3ol428k.

[251] Siemens. SIMATIC S7-1200 CPU 1212C - Data Sheet. https://sie.ag/3MSWnIX, 2022. Last Accessed: 16-03-2022.

[252] Johannes I. Single, Jürgen Schmidt, and Jens Denecke. State of research on the automation of HAZOP studies. *Journal of Loss Prevention in the Process Industries*, 62:103952, 2019. ISSN 0950-4230. doi: https://doi.org/10.1016/j.jlp.2019.103952. URL https://www.sciencedirect.com/science/article/pii/S0950423019302323.

[253] Sebastian Klovig Skelton. Destruction and Integrity Cyber Attacks on the Rise. https://bit.ly/3ByOrWn, 2021. Last Accessed: 01-11-21.

[254] Joe Slowik. Spyware stealer locker wiper: Lockergoga revisited. https://bit.ly/3neA1Eb, 2019. Last Accessed: 01-07-2020.

[255] Joseph Slowik. Evolution of ICS attacks and prospects for future disruptive events. https://bit.ly/2ABnTu7., 2020. Last Accessed: 01-07-2020.

[256] William Smart. Lessons Learned Review of the WannaCry Ransomware Cyber Attack. https://bit.ly/2UIg1AL, 2018.

[257] Paul Smith. The iran steel industry cyber attack explained. https://blog.scadafence.com/the-iran-steel-industry-cyber-attack-explained, 2022. Last Accessed: 27-03-2022.

[258] Paul Smith, Ewa Piatkowska, Edmund Widl, Filip Pröstl Andrén, and Thomas I Strasser. Towards a systematic approach for smart grid hazard analysis and experiment specification. In *2020 IEEE 18th International Conference on Industrial Informatics (INDIN)*, volume 1, pages 333–339. IEEE, 2020.

[259] Jae-Gu Song, Jung-Woon Lee, Cheol-Kwon Lee, Kee-Choon Kwon, and Dong-Young Lee. A Cyber Security Risk Assessment for the Design of I&C Systems in Nuclear Power Plants. *Nuclear Engineering and Technology*, 44, 12 2012. doi: 10.5516/NET.04.2011.065.

[260] Murugiah Souppaya and Karen Scarfone. NIST Special Publication 800-83: Guide to Malware Incident Prevention and Handling for Desktops and Laptops. Technical report, National Institute of Standards and Technology, 2013.

[261] Mike Spisak and James Darwin. Network Security Architecture. https://ibm.co/3hBztHl. Last Accessed: 30-06-2021.

[262] Alexander Staves. ICS/OT testing data sets and scripts. https://github.com/Warschak/ICS-OT-testing-data-sets-and-scripts, 2022.

[263] Alexander Staves, Harry Balderstone, Benjamin Green, Antonios Gouglidis, and David Hutchison. A framework to support ICS cyber incident response and recovery. In *the 17th International Conference on Information Systems for Crisis Response and Management, ISCRAM 2020 ; Conference date: 24-05-2020 Through 27-05-2020*, 2020. URL https://www.drrm.fralinlifesci.vt.edu/iscram2020/index.php.

[264] Alexander Staves, Tom Anderson, Harry Balderstone, Benjamin Green, Antonios Gouglidis, and David Hutchison. A Cyber Incident Response and Recovery Framework to Support Operators of Industrial Control Systems. *International Journal of Critical Infrastructure Protection*, 2022. ISSN 1874-5482. doi: https://doi.org/10.1016/j.ijcip.2021.100505. URL https://www.sciencedirect.com/science/article/pii/S187454822100086X.

[265] Alexander Staves, Antonios Gouglidis, and David Hutchison. An Analysis of Adversary-Centric Security Testing within Information and Operational Technology Environments. *Digital Threats: Research and Practice*, 2022.

[266] Keith Stougger, Victoria Pillitteri, Suzanne Lightman, Marshall Abrams, and Adam Hahn. NIST Special Publication 800-82: Guide to Industrial Control Systems (ICS) Security, Revision 2. Technical report, National Institute of Standards and Technology, 2015.

[267] Craig Sweigart. SCORE Security Checklist. Technical report, SANS Institute, 2003.

[268] Tenable. Nessus Vulnerability Scanner. https://bit.ly/399yave, 2021. Last Accessed: 16-09-2021.

[269] The European Parliament and Council. Regulation (EU) 2016/679. https://bit.ly/3bsAIpB, 2016. Last Accessed: 01-11-21.

[270] The European Parliament and the Council of the European Union. Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union, amending Regulation (EU) No 910/2014 and Directive (EU) 2018/1972, and repealing Directive (EU) 2016/1148 (NIS 2 Directive). https://www.nis-2-directive.com/NIS_2_Directive_Articles.html, 2022.

[271] The Guardian. Petya cyber-attack: Cadbury factory hit as ransomware spreads to australian businesses. https://bit.ly/319AQoG, 2017. Last Accessed: 23-03-2020.

[272] Joe Tidy. How a ransomware attack cost one firm £45m. https://bbc.in/3eOnfZw, 2019. Last Accessed: 01-07-2020.

[273] Joe Tidy. British Airways Fined 20m Pounds Sterling over Data Breach. https://bbc.in/3pRTejv, 2020. Last Accessed: 01-11-21.

[274] Tridium. About tridium: Tridium's history. https://bit.ly/3arDCYN, 2015. Last Accessed: 01-07-2020.

[275] Sharif Ullah, Sachin Shelly, Amin Hassanzadeh, Anup Nayak, and Kamrul Hasan. On the effectiveness of intrusion response systems against persistent threats. In *2020 International Conference on Computing, Networking and Communications (ICNC)*, pages 415–421. IEEE, 2020.

[276] United States Government Accountability Office. Cybersecurity Guidance is Available, but More Can Be Done to Promote Its Use. https://www.gao.gov/assets/gao-12-92.pdf, 2011.

[277] David I Urbina, Jairo A Giraldo, Alvaro A Cardenas, Nils Ole Tippenhauer, Junia Valente, Mustafa Faisal, Justin Ruths, Richard Candell, and Henrik Sandberg. Limiting the impact of stealthy attacks on industrial control systems. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 1092–1105, 2016.

[278] U.S Government Publishing Office. S.1353 - Cybersecurity Enhancement Act of 2014. https://bit.ly/35JFUm3, 2014.

[279] Pieter Van Vliet, M-T Kechadi, and Nhien-An Le-Khac. Forensics in industrial control system: a case study. In *Security of Industrial Control Systems and Cyber Physical Systems*, pages 147–156. Springer, 2015.

[280] Don C. Weber and Just Searle. Industrial Protocols Cheat Sheet v1.0. https://bit.ly/3IAHPKM, 2021. Last Accessed: 08-07-2021.

[281] Jody R. Westby. Governance of Enterprise Security: CyLab 2021 Report. https://bit.ly/3BdDefs, 2012.

[282] Westermo. Managed Ethernet Switch. https://www.westermo.com/products/ethernet-switches/layer-2/l110-f2g, 2019.

[283] Steve Whitty and Tony Foord. Is HAZOP worth all the effort it takes? https://www.icheme.org/media/9650/xxi-paper-022.pdf, 2009. Last Accessed: 20-08-2023.

[284] Tina Wu, Jules Ferdinand Pagna Disso, Kevin Jones, and Adrian Campos. Towards a SCADA forensics architecture. In *1st International Symposium for ICS & SCADA Cyber Security Research 2013 (ICS-CSR 2013) 1*, pages 12–21, 2013.

[285] Weider D. Yu, Dipti Baheti, and Jeremy Wai. Real-Time Operating System Security. https://bit.ly/3vZ3m9r.

[286] Kim Zetter. Meet 'flame,' the massive spy malware infiltrating iranian computers. https://bit.ly/2RS8vz3, 2012. Last Accessed: 01-07-2020.

[287] Kim Zetter. Everything we know about ukraine's power plant hack. https://bit.ly/3eEnDJV, 2016. Last Accessed: 02-07-2020.

[288] Christina Zhao. SolarWinds, Probably Hacked by Russia, Serves White House, Pentagon, NASA. https://bit.ly/3FB5U3H, 2020. Last Accessed: 11-10-2021.

[289] Bonnie Zhu, Anthony Joseph, and Shankar Sastry. A taxonomy of cyber attacks on SCADA systems. In *2011 International conference on internet of things and 4th international conference on cyber, physical and social computing*, pages 380–388. IEEE, 2011.

Appendix A

Standards And Guidelines Study Interview Guide

Pre-Screening

The purpose of pre-screening is to establish the validity of participants within the use of standards and guidelines for ICS/OT cyber incident response and recovery. This ensures that participants have the relevant background for discussing activities conducted for cyber incident response and recovery as well as the implementation of standards and guidelines to assess and improve these. As a starting point, the target interview audience will be composed of individuals that can be involved in cyber incident response and recovery activities for ICS/OT environments.

- ICS/OT Cyber Security Personnel
- ICS/OT System Operators
- ICS/OT System Managers
- ICS/OT System Engineers (Instrumentation, Control, and Automation)
- Operational Safety Managers
- ICS/OT Regulators

In selecting an appropriate participant sample, the aim is to understand the topic area from all relevant perspectives. To achieve this, a broad approach to the targeting of participants was applied. This resulted in a diverse collection of role-profiles. More specifically, those engaging in cyber incident response and recovery processes across multiple systems, with varying

levels of responsibility. This sampling approach provides multiple perspectives, building a broader picture of how cyber incident response and recovery activities are conducted.

Preface

The following question-set and associated notes will be applied during the preface phase.

- Reiterate the purpose of the interviews based on the interview guide, and the expected time-scale.
- Confirm the participant knows the full interview will be recorded, and that they will be told when the recording is due to begin, and when it is due to end.
- Turn ON the recording now.
- Confirm the consent of participation, that recorded has begun, and their rights with regards to participation.

Establishing Demographics

The following question-set and associated notes will be applied to the demographics phase.

- Please can you tell us your job title, and provide a brief overview of your core roles and responsibilities?
  Probe: Ask for clarity on any terms that are not clear.
- How many years of experience do you have working in this role?
  Probe: How many years of experience do you have working in this field?
  Probe: How many years of experience do you have working in this sector?
- At a very high level, please can you explain to us what you understand the term Response and Recovery to mean within the context of an Operational Technology (Industrial Control Systems) cyber security incident?
  Definition: Decisions and actions for the rapid implementation of a coordinated, multidisciplinary process, to manage the direct effects of an incident through protection of operational systems, human life, and the environment, creating the conditions required for a return of service.
  Definition: The process of rebuilding and restoring services to normal operation following an incident. Although distinct from response, recovery forms an integral part of response processes, as actions taken during the response activities can influence longer-term outcomes.

Scenario Familiarisation

The following question-set and associated notes will be applied during the scenario familiarisation phase.

- Please review the following infrastructure diagram (see Figure 3.1, a description will also be provided).
  Probe: Are there any aspects of the diagram which are unclear, or that you would like additional information on?
- Please review the following cyber incident diagram (see Figure 3.2, a description will also be provided).
  Probe: Are there any aspects of the diagram or attack which are unclear, or that you would like additional information on?

Response and Recovery Analysis

The following question-set and associated notes will be applied during the response and recovery analysis phase.

- Given your role in the organisation, at a high level, what are the core steps you would go through as part of response and recovery operations in the example scenario?
  Probe: Explore unusual terms and elaborate on anything that is unclear.
  Probe: Explore identified phases/processes.
  Probe: Is there anything unusual in this scenario that would cause you to deviate from a standard response process?
- How many individuals within the organisation would work directly with you on these steps, i.e., performing the same role as you or under your management?
- Who else would you have direct engagement with during response and recovery operations?
- How many individuals across the organisation would be involved in response and recover operations more generally speaking?
  Probe: Explore the use of any third-parties.
- When undertaking a response and recovery operation to this scenario, what do you consider the primary goal to be?
- When you are undertaking individual response and recovery actions, how do you factor in risk evaluation as part of the decision-making process?

Definition: Evaluating risk associated with the execution of specific actions, and thus the potential for unintended consequences arising as a result of those actions.

- Typically, what are the expected outputs post incident? So, once you have appropriately recovered from an incident and everything is back to normal?
  Definition: Reporting internally/externally, documenting, etc.
  Probe: Explore unknown/unclear post-incident outputs.
- Please review this second cyber incident diagram (see Figure 3.3, a description will also be provided). Would anything be done differently compared to the first scenario?
  Probe: Are there any aspects of the diagram or attack which are unclear, or that you would like additional information on?

Guidance Analysis

The following question-set and associated notes will be applied during the external guidance analysis phase.

- In your opinion, which standards or guidelines best cover response and recovery in relation to Operational Technology cyber-attacks targeting the nuclear sector?
  Probe: Why effective/not effective?
- As a final question, what is your opinion on currently available standards and guidelines within the context of cyber incident response and recovery?
  Probe: Why effective/not effective?

Conclude

The following question-set and associated notes will be applied during the conclusion phase.

- Confirm that the interview questions have been completed, and ask the interviewee if they would like to add anything in addition which may be relevant.
- If supporting documentation has been described and offered throughout the process, politely remind the interviewee to forward it on via E-Mail.
- Turn OFF the recording now.
- Thank the interviewee for their time and input into the project.
- Inform the interviewee that if at any time they recall any additional points deemed relevant to the discussed topic area, that one would greatly appreciate them being sent via E-Mail.
- Reiterate the options for withdrawal as described in the participant information sheet.

Appendix B

Scoping Framework Evaluative Study Interview Guide

Pre-Screening

The purpose of pre-screening is to establish the validity of participants within the evaluation of the framework. This ensures that participants have the relevant background for understanding the framework operation and providing feedback on its effectiveness and useability in practice. As a starting point, the target interview audience will be composed of individuals that can be involved in the design, planning or operation of an adversary-centric security test within ICS/OT environments. The roles included for participation in the evaluation of the framework are as follows:

- ICS/OT Cyber Security Personnel.
- ICS/OT Cyber Security Managers.
- ICS/OT Cyber Security Consultants.
- ICS/OT Penetration Testers.
- ICS/OT System Engineers.
- ICS/OT System Managers.

Reasonable endeavours will be made to ensure that there is appropriate diversity and representation within the interviewed sample group. However, it is noted that there exists well-known diversity issues within both the ICS/OT and cyber security community.

Preface

The following question-set and associated notes will be applied during the preface phase.

- *Reiterate the purpose of the interview based on the interview guide.*
- *Confirm that the participant knows that the full interview will be recorded and that they will be told when the recording begins and ends.*
- *Turn ON the recording now.*
- *Confirm the interviewee's consent to participation, that the recording has begun, and the rights of the interviewee in regards to their participation.*

Establishing Demographics

The following question-set and associated notes will be applied during the demographics phase.

- Please can you tell us your job title and provide a brief overview of your core roles and responsibilities? *Probe: Ask for clarity on any terms that are not clear.*
- How many years of experience do you have working in this role? *Probe: How many years of experience do you have working in this field? How many years of experience do you have working in this sector?*
- At a very high level, please can you explain to us what you understand the term adversary-centric security test to mean? *Note: Clarify with example engagements such as penetration test, vulnerability assessment or red team engagement if required. Definition: An adversary-centric security test is a form of engagement performed within a specified environment where an internal or external party is tasked with emulating techniques used by threat actors to identify existing vulnerabilities, assess existing defensive capabilities and/or train and evaluate incident response capabilities as part of the overall cyber risk management process.*
- Have you ever been involved in an adversary-centric security test that was performed for an ICS/OT environment? *Probe: If not, have you ever been involved in an adversary-centric security test in general?*
- At a very high level, what are the greatest challenges of conducting adversary-centric security tests for ICS/OT environments? *Probe: How does the operational nature of ICS/OT environments affect adversary-centric security testing?*

Framework Familiarisation

The following question-set and associated notes will be applied during the framework familiarisation phase.

- *Provide a detailed explanation of the framework and its different phases: TiDICS Layer Selection, Hazard Identification, Hazard Decomposition, Risk Quantification, Incorporation into the Overall Scoping Methodology.*
- Are there any aspects of the framework which are unclear or that you would like additional information on?

Framework Evaluation (no Application Scenario)

The following question-set and associated notes will be applied during the framework evaluation phase prior to its application in a scenario.

- What is your opinion on using the TiDICS model for separation of testing zones and layers based on safety and operational risks? *Probe: Do you think the model is appropriate for this?*
- Do you agree that the types of risk that adversary-centric security testing presents to safety and the operational process are comprehensively considered within the framework? *Probe: What other types of risk would you expect to identify?*
- What challenges could affect the collection or quality of data for risk quantification? *Probe: How could these challenges be overcome?*
- From the framework's overview, do you think the output of the framework could be used in the overall scoping of an adversary-centric security test? *Probe: Why? Why not?*
- From the framework's overview, do you believe that the output of the framework is accurate enough to ensure a full understanding of the safety and operational risks from adversary-centric security testing on ICS/OT so that depth of testing can be maximised while minimising risk to the operational process? *Probe: Why? Why not?*
- From the framework's overview, do you believe that the framework can be applied in all ICS/OT environments where safety and operational risks are a concern? *Probe: For which environments would it not be applicable?*

Scenario Familiarisation

The following question-set and associated notes will be applied during the application scenario of the framework familiarisation phase.

- *Provide a detailed explanation of the application of the framework on an example scenario (see figure 6.2 and the different methodologies used in each phase of the framework: TiDICS Layer Selection, Hazard Identification with (C)HAZOP, Hazard decomposition with FTA, Risk quantification of specific tools or techniques (Nmap, Nessus etc), Incorporation into the Overall Scoping Methodology.*
- Were you familiar with using (C)HAZOP to identify hazards prior to its use within this scenario? *Probe: Have you ever used (C)HAZOP before? In what context?*
- Were you familiar with using FTA for decomposing hazards into smaller basic events prior to its use within this scenario? *Probe: Have you ever used FTA before? In what contextowever, despite the framework being supported by data collected in an ICS/OT testbed composed of real-world hardware used in industrial environments, further evaluation was required to determine the framework's validity.?*

Framework Evaluation (with Application Scenario)

The following question-set and associated notes will be applied during the re-evaluation of the framework phase when applied within an example scenario.

- Does your opinion of the framework's validity, accuracy, and applicability change when presented with an example application of its use? *Probe: Why? Why Not?*
- What is your opinion on the use of (C)HAZOP for identifying hazards that could occur during an adversary-centric security test on ICS/OT? *Probe: Are there any other methodologies for identifying hazards that you would prefer to use over (C)HAZOP?*
- What is your opinion on the use of FTA to decompose hazards into smaller basic events for use in risk quantification? *Probe: Are there any other methodologies for decomposing hazards that you would prefer to use over FTA?*
- What is your opinion on the methodologies used for quantifying the risk of basic events?
- Overall, would you use this framework as part of the overall scoping methodology for an adversary-centric security test on ICS/OT? *Probe: Why? Why Not?*

Conclude

The following question-set and associated notes will be applied during the conclusion phase.

- *Confirm that the interview questions have been completed and ask the interviewee if they would like to add anything in addition which may be relevant.*
- *If supporting documentation has been described and offered throughout the process, politely remind the interviewee to forward it on via E-Mail.*
- *turn OFF the recording now.*
- *Inform the interviewee that if at any time they recall any additional points relevant to the discussed topic area, that one would greatly appreciate them being sent via E-Mail.*
- *Reiterate the options for withdrawal as described in the participant information sheet.*