HORIZON EUROPE PROGRAMME
HORIZON-CL3-2021-CS-01-01

A EUROPEAN CYBER RESILIENCE FRAMEWORK WITH ARTIFICIAL INTELLIGENCE -ASSISTED ORCHESTRATION & AUTOMATION FOR BUSINESS CONTINUITY, INCIDENT RESPONSE & INFORMATION EXCHANGE

D4.2: Coordinated Response & Preparedness Enablers v2

Abstract: This document reports on the second and final release of enablers developed under PHOENI2X WP4 ("Coordinated Response & Preparedness Enablers"), specifically covering the outputs of Task 4.1 ("Resilience Orchestration, Automation & Response"), Task 4.2 ("Resilience Playbooks Specification, Translation & Lifecycle Management"), Task 4.3 ("Cyber Range & Serious Games for Resilience Assessment & Training"), and Task 4.4 ("Alerting, Reporting & Information Exchange"). The main goal of this document is to accompany the release of the enablers, providing details on their design, development and usage.

Contractual Date of Delivery	31/05/2025
Actual Date of Delivery	02/06/2025
Deliverable Security Class	PU
Editor	Mateusz Zych (UiO)
Contributors	UPAT, SANL, COSM, FGC, PPC, WSE, AEGIS, SEA, ATOS, APS, NPS, UiO, DSA
Quality Assurance	SEA, NCSA

This project has received funding from the Horizon Europe Research and Innovation programme under Grant Agreement No 101070586

Document Revisions & Quality Assurance

Internal Reviewers

Review 1: SEA

Review 2: NCSA

Revisions

Version	Date	By	Overview
0.1	14/04/2025	Editor	Initial draft, with ToC & assignments
0.2	30/04/2025	ATOS	SMIR updates
0.3	19/05/2025	SANL	Resilience Orchestration, Automation & Response Resilience Playbooks Specification, Translation & Lifecycle Management The Resilience Cyber Range
0.4	23/05/2025	SEA	Serious Games
0.5	23/05/2025	UiO	Introduction, Notification Playbook, Playbook exchange, Conclusion
0.6	23/05/2025	UiO	Formating
0.7	27/05/2025	SEA	#1 Review
0.8	30/05/2025	NCSA	#2 Review
0.9	31/05/2025	UiO	Addressed the reviews

Contents

List of Tables

List of Figures

List of Abbreviations

AI: Artificial Intelligence

APT: Advanced Persistence Threat

ATM: Attack Categorization Modeling

BC: Business Continuity

CACAO: Collaborative Automated Course of Action Operations

CI: Critical Infrastructures

CNN: Convolutional Neural Network

CRC: Cyber Resilience Center

CTI: Cyber Threat Intelligence

DoA: Description of Action

DTL-EL: Deep Transfer Learning for Exploit Labelling

DX.X: Deliverable X.X

ENISA: European Union Agency for Cybersecurity

EU: European Union

FAIR: Factor analysis of Information Risk

FSM: Finite State Machine

GAN: Generative Adversarial Network

GUI: Graphical User Interface

IoC: Indicator of Compromise

IR: Incident Response

KPI: Key Performance Indicator

LSTM: Long Sort Term Memory

ML: Machine Learning

MS: Member State

MSA: Micro-Service Architecture

MSRA: Maritime Supply Chain Risk Assessment

MTTA: Mean Time To Acknowledge

MTTR: Mean Time To Remediate

MVP: Minimum Viable Product

NLP: Natural Language Processing

OES: Operators of Essential Services

PMEM: Predictive Maintenance

RCR: Resilience Cyber Range

ROAR: Resilience Automation, Orchestration, and Response

RP: Resilience Playbook

SIEM: Security Information & Event Management System

SOAR: Security Orchestration, Automation and Response

TIP: Threat Intelligence Platform

TL: Transfer Learning

TTP: Tactics, Technics and Processes

UEBA: User and Entity Behavior Analytics

WP: Work Package

XDR: eXtended Detection and Response

1 Introduction

1.1 Purpose of the Document

This deliverable is the second output of WP4 ("Coordinated Response & Preparedness Enablers"), documenting the delivery of the second and final version of the relevant components of PHOENI2X. As such, it documents the delivery of enablers coming from the underlying tasks, including:

- **Task 4.1 - Resilience Orchestration, Automation & Response**, providing components enabling to the automation & orchestration of resilience (incident response & business continuity) – related workflows. These results are documented in Section 2.
- **Task 4.2 - Resilience Playbooks Specification, Translation & Lifecycle Management**, providing the machine-processable & executable playbooks themselves that encode the above workflows. These results are documented in Section 3.
- **Task 4.3 - Cyber Range & Serious Games for Resilience Assessment & Training**, providing the components supporting the preparedness of involved users (both concerning testing & training at the process, but also the overall training and awareness at the human level). These results are documented in Section 4.
- **Task 4.4 - Alerting, Reporting & Information Exchange**, providing the components and processes enabling the timely production & exchange of information, as needed both operationally (e.g., to ensure timely alerting & response), but also from a regulatory perspective. These results are documented in Section 5.

Details on the above will be provided in the sections that follow, with each Task having a dedicated section, as noted. Finally, Section 6 provides the concluding remarks.

1.2 Methodology & Deliverable Mapping

Overall, the above WP4 components cover a big part of the overall PHOENI2X framework, as visualised in Figure 1.

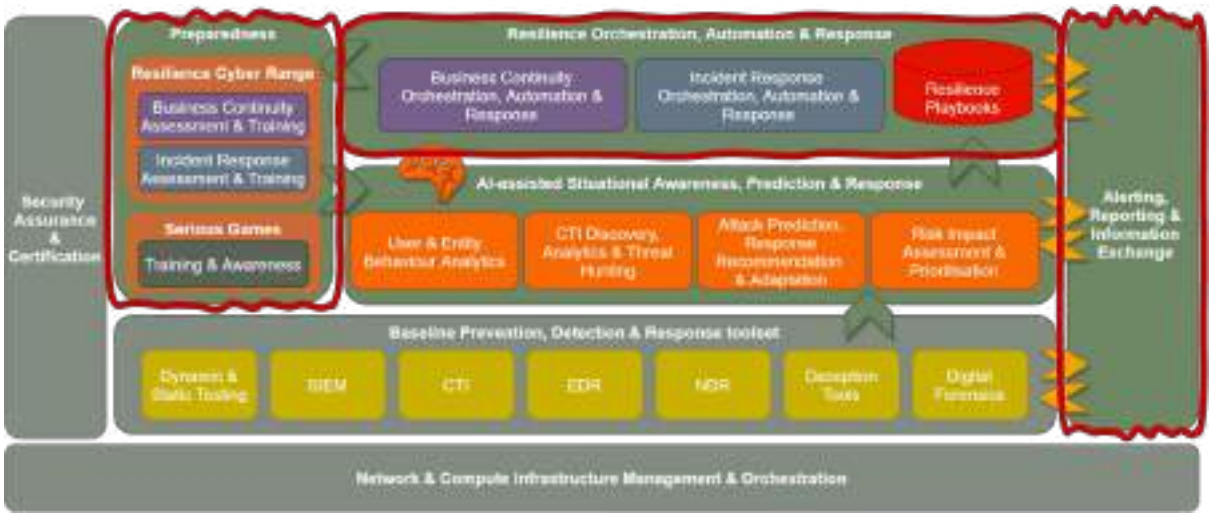

Figure 1: WP4 components within the PHOENI2X high-level architecture.

In fact, WP4 outputs provide key enablers aligned with (and necessary to achieve) 3 of the project's overarching objectives, including:

- **Objective 2** (*"To design & develop Resilience Orchestration, Automation and Response mechanisms, encompassing proactive and reactive business continuity, recovery and incident handling tasks"*), covered by Tasks 4.1 & 4.2.
- **Objective 3** (*"To offer enhanced Preparedness through a Resilience Cyber Range and Serious Games"*), covered by Task 4.3.
- **Objective 4** (*"To provide Alerting, Reporting & Information Exchange mechanisms & processes enabling collaboration between private and public critical sector actors at the national and European level"*), covered by Task 4.4.

From a practical perspective, the component design & development activities described herein used the outputs of WP2 as foundational inputs, including the component-level specifications, functional & non-functional requirements and interfacing requirements stemming from the PHOENI2X architecture, as documented in D2.1 ("PHOENi²X Requirements & Architecture"). With the above at hand, the development of the individual WP4 components proceeded, with input from other activities & actors (e.g., use case owners), as needed. These initial developments were described and delivered in D4.1 - "Coordinated Response & Preparedness Enablers v1". Deliverable D4.2 builds on top of D4.1 and provides the status updates and developments performed during the reporting period.

The natural outcome of WP4 outputs (along with the outputs of WP3) is WP5, for the integration and eventual demonstration & validation of PHOENI2X. In this context, valuable input for the development of V1 of WP4 components was also provided by the early integration & demonstration efforts carried out within WP5 to derive the MVP version of PHOENI2X (as documented in D5.1 - "PHOENI2X framework - MVP") and the final release of the components reported in D5.3 - "PHOENI2X framework - Final".

2 Resilience Orchestration, Automation & Response

2.1 Overview

The Resilience Orchestration, Automation & Response (ROAR) component of PHOENI2X is developed by extending the SPHYNX Incident Response (IR) tool. This system facilitates both manual and automated execution of Collaborative Automated Course of Action Operations (CACAO [1]) security playbooks (for more information on playbooks, refer to Section 3), following the latest specification. ROAR offers a graphical drag-and-drop interface for designing and editing CACAO playbooks, which can later be executed directly or exported as CACAO-compliant JSON files.

In essence, this component functions as a Security Orchestration, Automation, and Response (SOAR) solution, supporting the full lifecycle of cybersecurity operations from prevention and detection to investigation and response.

Finally, the ROAR component can operate as a standalone solution or be integrated as a module within the SPHYNX Security and Privacy Assurance Suite (SPA), as described in Section 4 of Deliverable D2.2 “Baseline Enablers”.

For more information about the key features of this component and how it is positioned within the high-level and detailed architecture of PHOENI2X, please refer to Section 2.1 of Deliverable D4.1.

2.2 Design Details

The ROAR tool includes an interactive dashboard that enables users to specify, edit, and execute playbooks. It also offers real-time insights into the system’s status and playbook execution, displaying both high- and low-level logs, key performance indicators (KPIs), user notifications, and other relevant information. During the project’s lifecycle, ROAR's editor has been expanded to fully support the functionalities of each playbook step defined at the latest CACAO specification. Also, playbook management functionalities, such as "Save", “Save as” and “Load” were implemented into the editor. An example of how a security playbook that follows the CACAO v2 specification can be defined within the ROAR GUI is presented in Figure 2 where the security ticket creation playbook used in Use Case 2 is being designed.

Overall, as the ROAR design & capabilities are fully tailored to support and implement the CACAO v2 specification, we defer the reader to Section 3 below, where we provide more details on the Resilience Playbooks themselves (e.g., on the design of each of the playbook steps supported by ROAR).

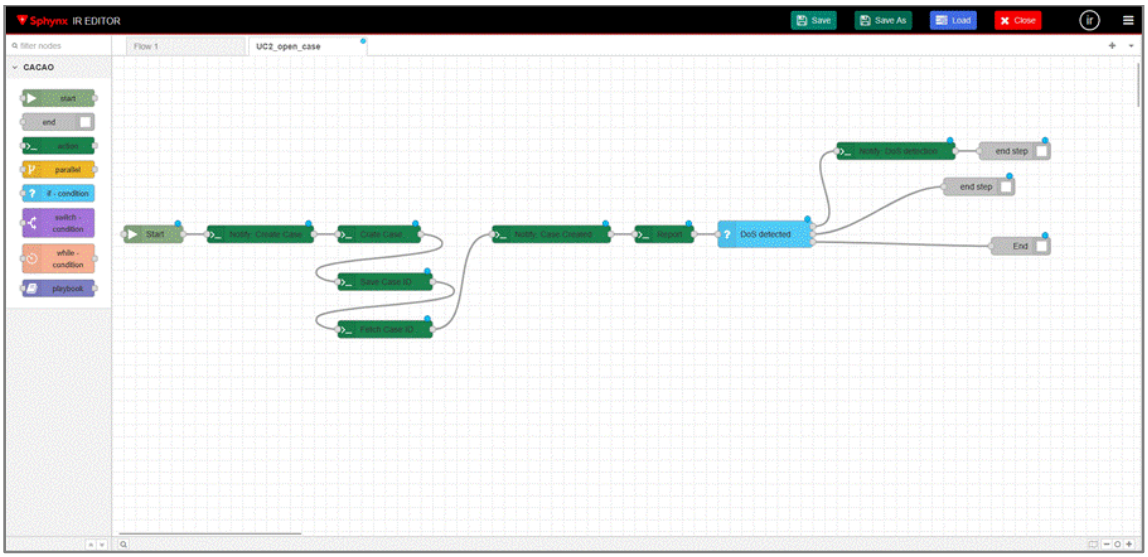

Figure 2: Updated ROAR CACAO playbook editor. Design of security ticket generation playbook utilized in Use Case 2.

2.3 Implementation Details

To support the needs of PHOENI2X pilots, ROAR's internal architecture was extended to support message brokers, such as Kafka, as well as custom execution agents to operate on SSH and HTTP-API targets, as illustrated in Figure 3.

The components found in the ROAR tool can be deployed either on a single host or operate distributed across multiple hosts, depending on the deployment needs. Each component runs in its own isolated Docker container, although native (non-containerized) deployment is also supported but not recommended. For optimal security and reliability, the components should ideally be deployed on a secure and monitored infrastructure, external to the target organization, to ensure continuous, tamper-resistant operation.

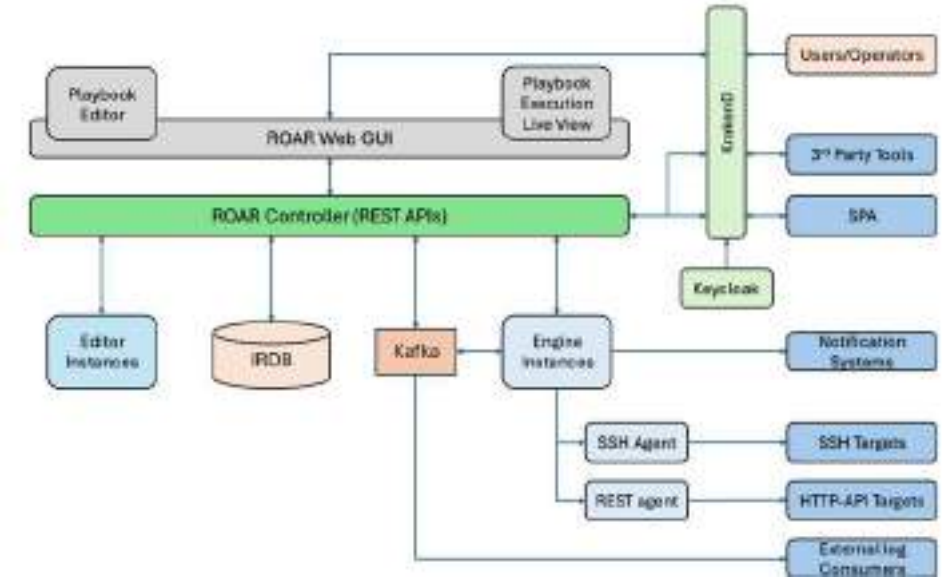

Figure 3: ROAR final architecture.

For more information about the ROAR components shown in Figure 5 (Web GUI, Controller, Engine, Editor, and IRDB), as well as a detailed view of ROAR's internal architecture, its integration with external components, and an overview of the process for importing and exporting Resilience Playbooks (RPs) in the CACAO format using the ROAR tool, please refer to Section 2.3 of Deliverable D4.1. During this project cycle the ROAR ecosystem was expanded to include a Kafka broker that streamlines the communication between the ROAR controller and the playbook execution engines. Moreover, ROAR utilizes Kafka to propagate logs and execution information that can also be consumed by 3rd party tools. Also, ROAR's latest implementation allows the production of logs to various simultaneous topics to serve multiple 3rd party tools concurrently. Security engineers can define each playbook's output topic when playbooks are loaded into the engine(s), enabling fine-grained log output.

Further, additional interfacing and reporting capabilities have been incorporated into ROAR to enable seamless integration with other PHOENI2X components. These enhancements support not only the execution of playbook actions but also improve PHOENI2X front-end interoperability. The integration between ROAR and the FVT component is illustrated in Figures 4 and 5 below. The elements shown in Figure 4 are implemented during the first cycle of the project. Figure 5 presents the components implemented in the second cycle, where playbooks are categorized as either IR or BC playbooks. It also displays the success rate of their execution, or the failure rate if any issues occurred.

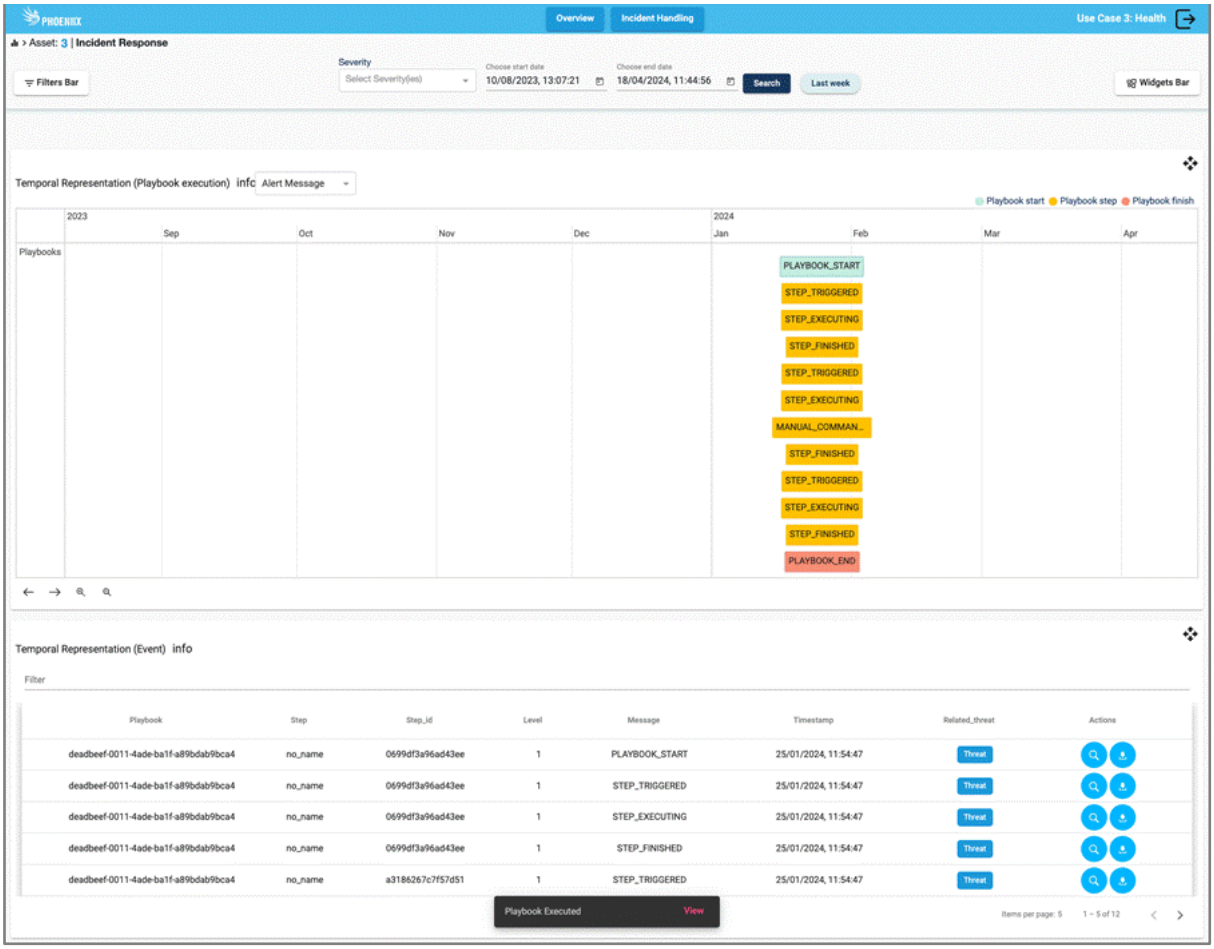

Figure 4: Integration of ROAR actions onto the FVT.

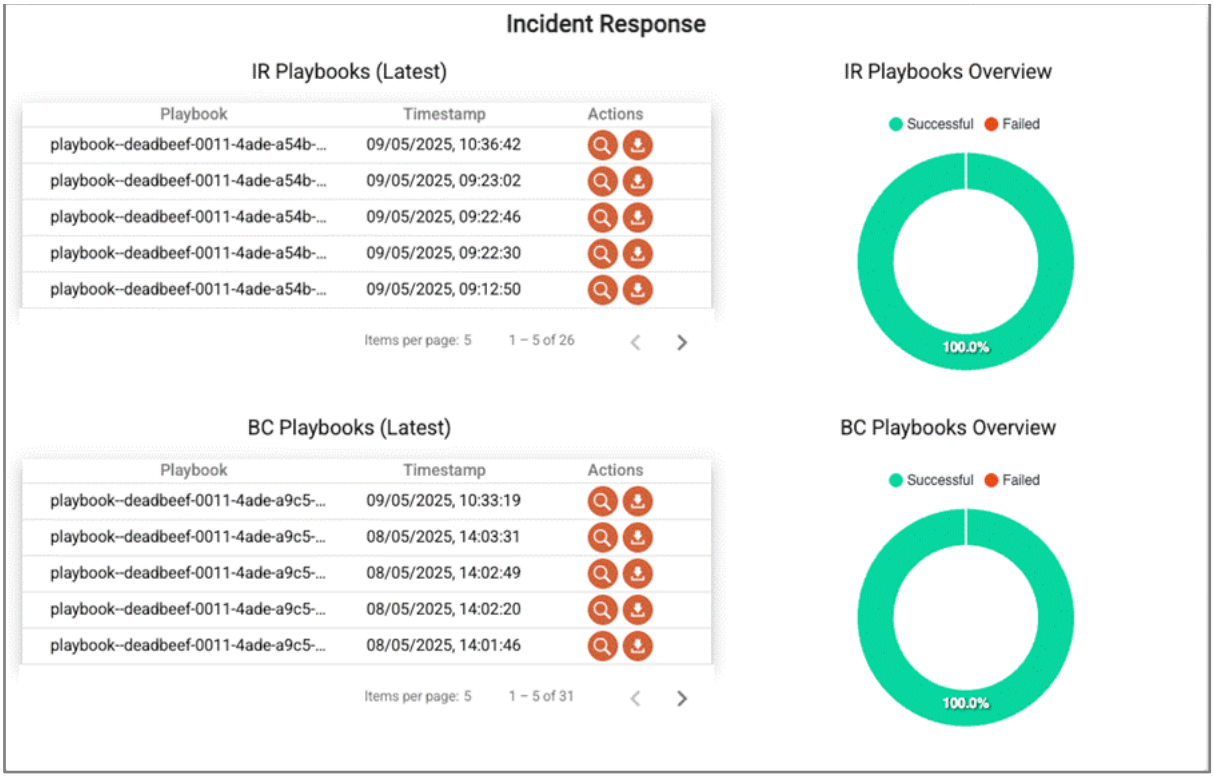

Figure 5: Executed Playbooks Per Type & Success Rate.

2.4 Final Status

The focus on the second and final cycle of the project is shifted on extending the support for additional Incident Response (IR) Playbooks, including the novel Business Continuity (BC) Playbooks introduced in PHOENI2X. As a proof of concept, these various types of playbooks were designed, implemented, and successfully integrated into the three distinct use case scenarios of the project, covering different categories of Operators of Essential Services (OES): Energy, Transport, and Health.

In the Energy sector use case (UC1), ROAR was able to respond to Distributed Denial of Service (DDoS) attacks and physical attacks targeting the primary communication channel of smart meters.

In the Transport sector use case scenario (UC2), ROAR was configured by security professionals to respond to Data Tampering, Denial of Service (DoS) attacks, and Privilege Escalation attempts. Additionally, a Business Continuity (BC) playbook was executed to support the investigation process by collecting relevant logs and documentation.

In the Healthcare sector use scenario (UC3), ROAR was able to respond to incidents such as unauthorized commits in GitLab, unauthorized deletion of code branches, and the introduction of vulnerable libraries into the codebase. Furthermore, ROAR supported the execution of BC playbooks following these attacks, as well as reporting and alerting playbooks to notify relevant stakeholders.

As a result of these efforts, this final phase enabled broader integration with additional tools. The design and configuration of the playbooks were streamlined through the ROAR tool, which has undergone UI and architectural enhancements to better support the needs of security professionals.

Moreover, final integration with the FVT component has been completed to ensure that essential information such as executed playbooks and their types is now presented in the frontend of the PHOENI2X platform.

Additionally, the integration between ROAR and the UEBA module was developed. UEBA's predictive capabilities now allow it to detect abnormal behavior and notify the SPA Suite, which in turn triggers ROAR to execute the appropriate playbooks/actions. Finally, ROAR has also been integrated with the RCR. This integration allows for the assignment of targeted cybersecurity training to employees following an incident response event, reinforcing an organization's ability to learn from incidents and continuously improve its resilience posture.

2.5 Lessons Learnt and Recommendations for Future Work

Throughout the implementation and integration of the ROAR tool across the PHOENI2X use cases, several important lessons were learned.

The integration of ROAR with other PHOENI2X components such as the UEBA module, SPA Suite, and the FVT highlighted the need for clearly defined interfaces and synchronization mechanisms. Future efforts can be focused on adopting standardized integration frameworks to streamline development and reduce the risk of misalignment between components.

Looking ahead, several opportunities for further development have been identified. These include extending ROAR’s capabilities to support proactive threat hunting and threat intelligence enrichment as well as enhancing scalability for deployment in larger infrastructures. Additionally, deploying ROAR in live operational environments, instead of simulated ones, would provide valuable insights into its real-world performance and long-term operational readiness.

3 Resilience Playbooks Specification, Translation & Lifecycle Management

3.1 Overview

Resilience Playbooks (RPs), developed under WP4, Task 4.2, are a core element of the PHOENI2X framework and its ROAR capabilities. They offer a structured, machine-readable encoding for defining sequences of actions that support an organization's business continuity, recovery, and incident response activities. Each action reflects a fundamental task (e.g., adding a rule to a firewall). Thus, through RPs, organisations will be able to specify, automate the execution (via the purpose-built execution and orchestration engine), monitor the progress, and assess the effectiveness of all their business continuity, recovery, and incident response-related processes.

RPs are designed to be adaptable and contextualized via PHOENI2X's AI-driven situational awareness and post-incident insights, RPs are also customizable to each OES's regulatory and technical environment, shareable across organizations at machine speed, capable of supporting what-if analysis for planning purposes, and translatable for use in training and cyber range simulations.

For the placement of the playbooks and the associated databases in the overall PHOENI2X architecture please refer to section 3.1 of deliverable D4.1.

3.2 Design & Implementation Details

To support the above, PHOENI2X Resilience Playbooks (RPs) adopts and extends the OASIS Collaborative Automated Course of Action Operations (CACAO) standard, with version 2.0 of the specification officially approved in November 2023. This standard is a cybersecurity-specific schema and taxonomy for creating, documenting, and sharing playbooks in a structured and standardized format across organizational boundaries and technological solutions. For more information about the CACAO playbook structure and the various playbook types such as Business Continuity Playbooks and Incident Response Playbooks can be found in Section 3.2 of the Deliverable D4.1.

In short, the available types of playbooks, in relation to the incident occurrence timeline, are shown in Figure 6.

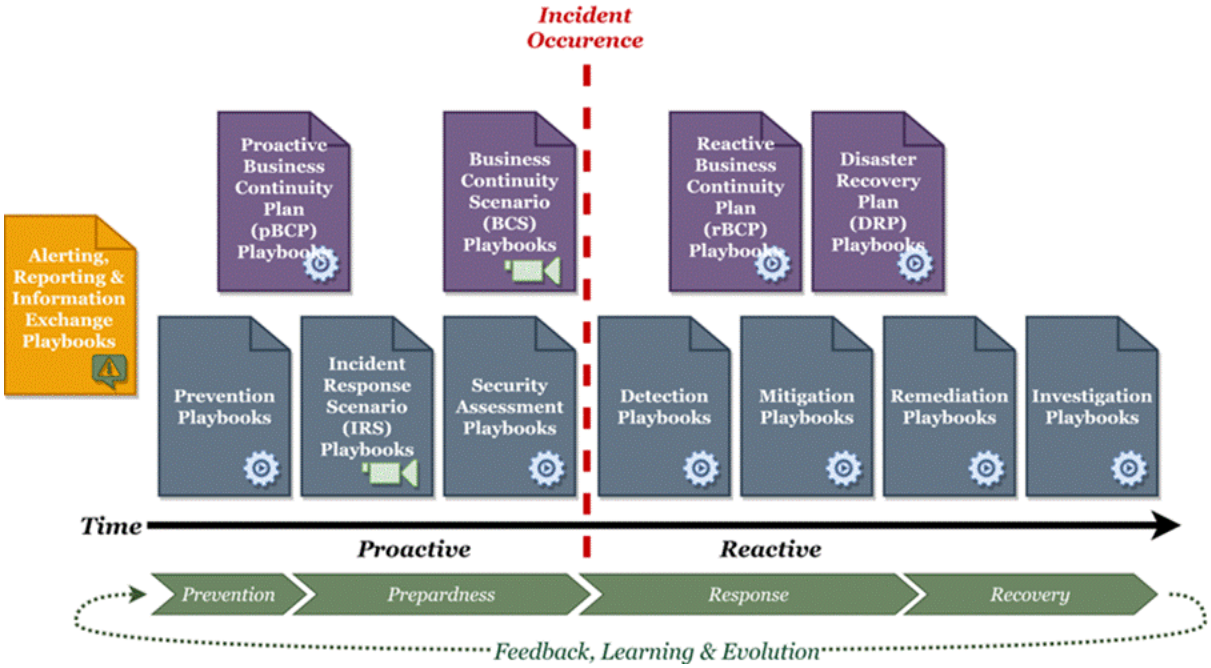

Figure 6: Different types of Resilience Playbooks envisioned in PHOENI2X [4].

In addition, it is worth mentioning that the consortium has developed and open-sourced JSON validation schemas for CACAO Version 2.0. The validation schemas are hosted at the official GitHub[1] of the OASIS CACAO technical committee and are considered the de facto validation mechanism for all implementations.

In PHOENI2X, Resilience Playbooks (RPs) are created, managed, and executed within the ROAR tool, which serves as the central platform for managing the entire lifecycle of RPs. This includes the initial design and specification of playbooks, their ongoing editing and refinement, as well as their automated or manual execution in response to incidents. ROAR provides a user-friendly, graphical interface that enables users to define complex workflows without requiring programming expertise.

For detailed information on the supported CACAO playbook steps (e.g., start step and end step), the playbook design process, and the definition of various step properties, please refer to Section 3.3 of Deliverable D4.1.

3.2.1 Supported CACAO Playbook Steps

ROAR's latest version was updated to support the design and execution of CACAO security playbooks using the step definitions described by the latest CACAO version. While the initial version of the ROAR tool supported the seven steps defined in previous versions of the specification, the current version supports all eight CACAO v2 steps, namely (i) start, (ii) end, (iii) action, (iv) parallel, (v) if-condition, (vi) switch-condition, (vii) while-condition, and (vii) playbook step. The list of available steps found in the latest ROAR version are presented in Figure 7. The following sub-sections present the development performed during this project cycle on each available step.

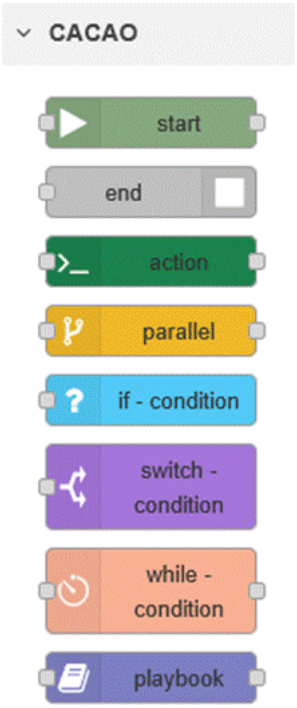

Figure 7: CACAO v2 playbook steps available as drag-and-drop enabled graph nodes within the ROAR GUI.

[1] https://github.com/oasis-open/cacao-json-schemas

3.2.2 Start Step

The playbook start step's editor was updated to present the playbook ID, creator, creation date and modification date in version 2 compliant formats. Also, ROAR handles the generation of these fields and users are prevented from tampering them. Moreover, calendar widgets are used to define the start and end dates of a playbook's validity period while ROAR generates the appropriate date and time formats.

Also, the updated editor allows the definition and editing of playbook variables supporting eleven data formats, namely (i) string, (ii) integer, (iii) long, (iv) UUID, (v) MAC_ADDR, (vi) IPV4_ADDR, (vii) IPV6_ADDR, (viii) URI, (ix) HEX string, (x) SHA256 hash, and (xi) dictionary. The variable editor performs data format validation when editing or adding variables. The updated start step editor is presented in Figure 8.

Moreover, ROAR extended the utilization of variables during playbook execution. The latest updates allow playbook steps to generate new variables during playbook execution so the results of each executed command can be saved and utilized by other steps. This allows for the creation of more complex playbooks without the need to predefine all the playbook variables prior. Also, ROAR introduces the functionality of exposing variable initialization and editing to the action steps. In this way, security engineers can utilize the action steps to alter the contents of variables via action steps, a feature that cannot be accomplished by the current CACAO specification.

Finally, the latest ROAR implementation enables log output to various Kafka topics so external tools can monitor playbook execution in an almost real-time fashion. To provide the highest possible degree of fine-grained monitoring, security engineers can utilize the start step's editor to direct copies of the playbook execution logs to various Kafka topics. This option can be enabled for an entire engine instance or in a per-playbook fashion, allowing playbooks co-residing in the same engine to provide copies of their logs each to a single or multiple different Kafka topics.

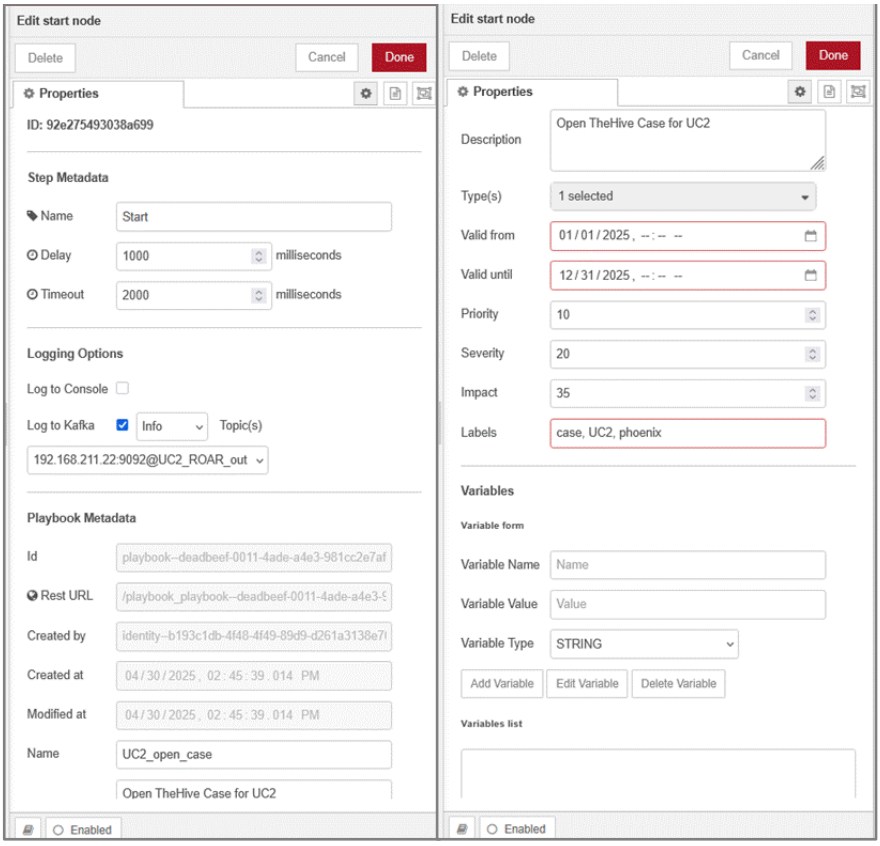

Figure 8: Updated Start step editor.

3.2.3 Action Step

The “single” step used in previous versions of the ROAR tool has been updated to the “action” step used in the latest version of the CACAO specification. During this process, we have updated the step’s editor as follows. First, Users can now select from a list of pre-configured targets where the step’s commands will operate on. This functionality eliminates human error and prevents users from designing destructive playbooks operating outside of the scope of the target infrastructure. The ROAR tool’s administrators can now define all the appropriate targets that should be available to the playbooks along with the required credentials for accessing them, such as SSH keys, passwords, tokens, etc. This process is performed using a dedicated editor with elevated privileges. During playbook development, the security engineers can select the appropriate targets by their ID while the sensitive credentials are hidden and cannot be altered during the playbook design process.

Moreover, the HTTP-API commands are extended to support all available HTTP operations, and the new action step menu provides visual editors for JSON fields, such as headers and bodies, providing visual formatting and input validation. The action step’s editor is presented in Figures 9 and 10.

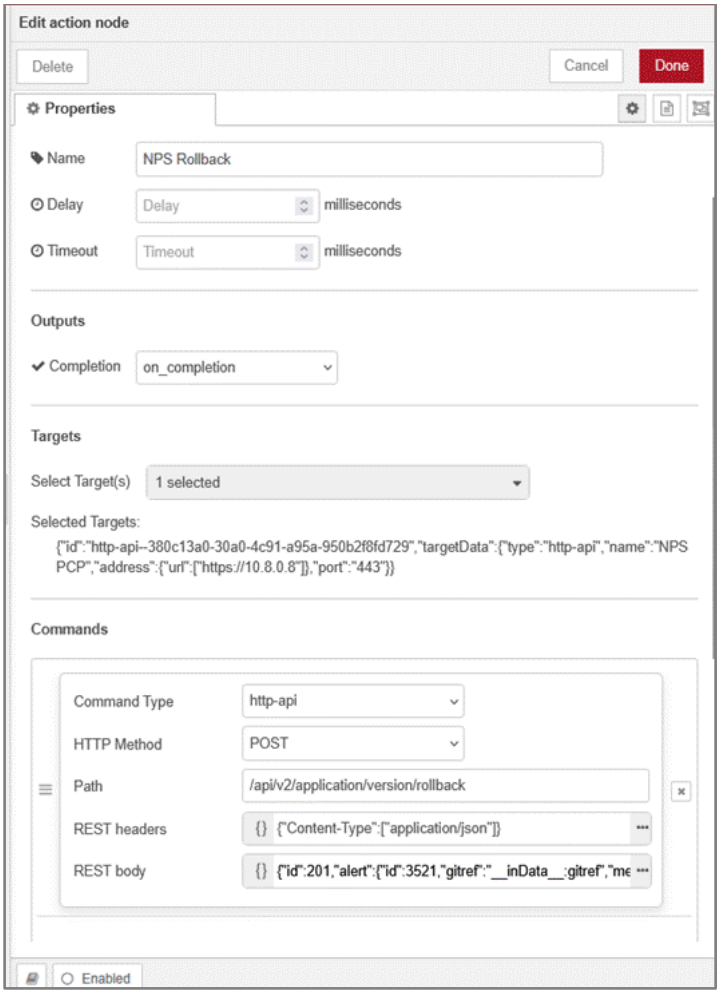

Figure 9: Action Step Editor.

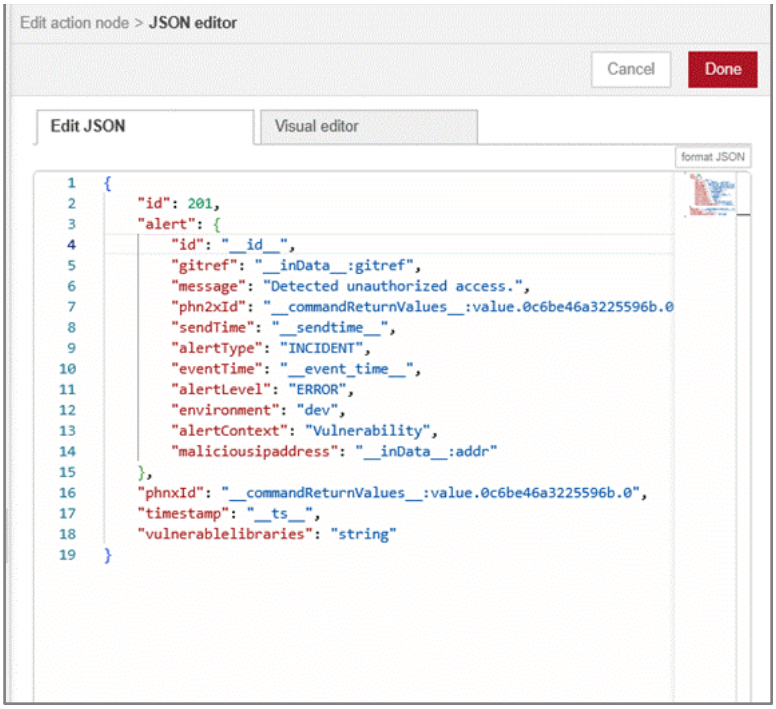

Figure 10: Action Step JSON field editor.

3.2.4 Playbook Step

The latest version of ROAR, developed during this project period introduces the playbook step defined in CACAO version 2. This step is utilized to call a target playbook from a caller playbook, serving as a "function call". This step can be utilized to daisy-chain multiple playbooks together and allows reusing previously defined playbooks without the need for re-implementing their functionality. This streamlines the playbook development process, allows the creation of more human-readable playbooks and enables quick changes and updates where users can edit a single target playbook utilized by several others. For example, a single playbook can be developed to create security incident tickets which can be utilized by various other playbooks. If the ticket creation API gets updated, the security engineers are required to edit only the ticket creation playbook.

ROAR also provides extended functionality to the playbook step, extending the CACAO definition, to support the tool's distributed nature as well as the pilot's needs. Via the editor, the security engineers can decide if the playbook step should wait for the target playbook to complete its operations before proceeding to the next step or proceed directly after calling it. This emulates synchronous and asynchronous function calling found in traditional programming languages. Also, since ROAR can support multiple concurrent execution engines (for isolation and scalability), the editor allows the specification of the desired engine where the callee playbook will be executed. The playbook step's editor is presented in Figure 11.

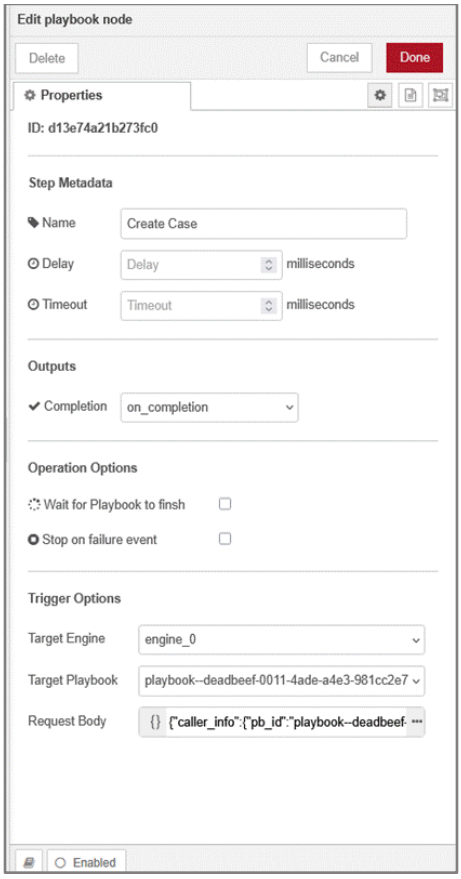

Figure 11: Playbook Step editor.

3.2.5 Conditional Steps

The CACAO playbook steps responsible for conditional logic, such as if-condition, switch-condition and while-condition steps were also updated according to the latest CACAO specification. In the latest ROAR version, the switch-condition step provides a default case, similar to the default case found in switch conditions provided by traditional programming languages. Also, all steps can operate on playbook variables or variables generated during playbook execution.

The most notable update performed on the conditional steps is the introduction of end steps indicating the end of logical branches. CACAO specification v2 dictates that every logical branch should be terminated with an end step. ROAR utilizes these end steps along with emulated call stacks where conditional steps save their ID before branching out. In this way, end steps are repurposed to enforce caller-callee relations between steps, similar to return values found in traditional programming languages. With this feature, security engineers do not need to specify the step to which the end of a branch should return to as it is now automatically handled by ROAR. An example of a playbook utilizing conditional steps with end steps for branch termination is presented in Figure 12.

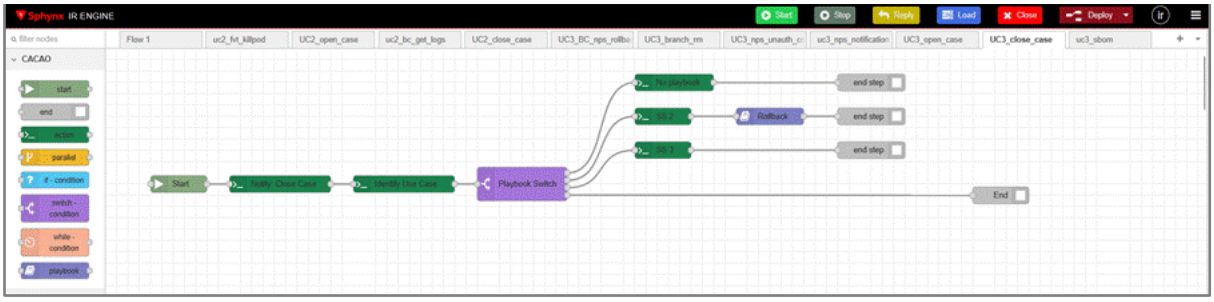

Figure 12: CACAO playbook used in Use Case 3 utilizing the switch condition step with branch termination using end steps.

3.3 Final Status

In the previous cycle of the project the emphasis was given on the delivery of the Incident Response Playbooks to support the use cases and the relevant demonstrators, with some preliminary work on the more novel application of playbooks for Business Continuity.

Specifically, the initial set of incident response playbooks covering the core categories of (i) Prevention, (ii) Security Assessment, (iii) Detection, (iv) Mitigation, (v) Remediation, and (vi) Investigation was developed during the first cycle. These playbooks were then tailored to align with the demonstration scenarios of each project use case (e.g., see Figure 13), and were subsequently demonstrated and validated as part of the evaluation of the first integrated version of the PHOENI2X framework.

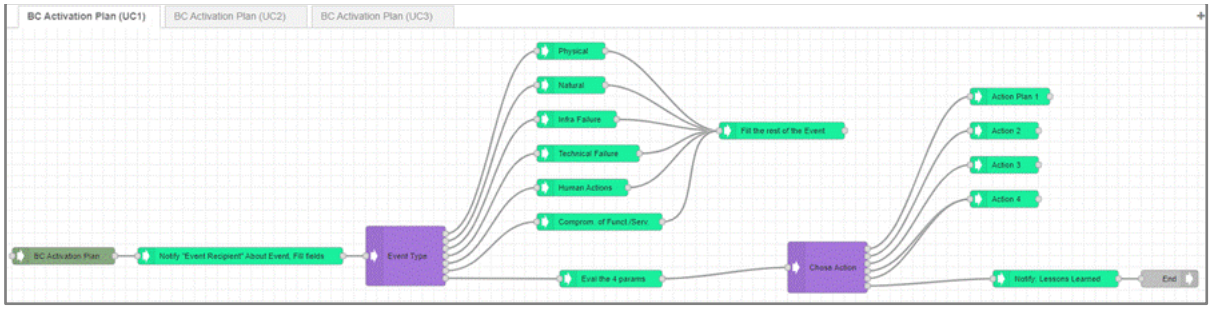

Figure 13: High-level BC process playbook [4].

In this second and final cycle of the project, the focus shifted to the design and delivery of additional IR Playbooks, as well as the delivery of the required BC-focused Playbooks and the Scenario Playbooks.

Use Case 1: Energy Sector

1. Detection Playbook, alternative #1
2. Detection Playbook, alternative #2
3. Business Continuity (BC) Playbook: SmartMeters Disconnection & Establishes connection through secondary WAN until main channel of communication is restored (Figure 14) *playbook added to the second cycle

Use Case 2: Transport Sector

1. Incident Response (IR) Playbook: Data Tampering Attack
2. Incident Response (IR) Playbook: DoS Attack (Figure 15)
3. Incident Response (IR) Playbook: Priviledge Escalation Attack
4. Business Continuity (BC) Playbook: Logs & Documentation *playbook added to the second cycle

Use Case 3: Healthcare Sector

1. Incident Response (IR) Investigation Playbook: Abnormal behaviour (unauthorized commit/push)
2. Incident Response (IR) Investigation Playbook: Abnormal behaviour (removal of code branch)
3. Incident Response (IR) Attack containment Playbook: Vulnerable library deferred detection (measures to contain attack)
4. BC Playbook: Rollback to version without the vulnerable library (Figure 16). *playbook added to the second cycle
5. Threat Ellimination Playbook: Create a new version without the vulnerable library
6. Reporting Playbook: Create incident in Hive
7. Alerting Playbook: Alerting other PHOENI2X instances

All the above playbooks in the three different use cases, were tailored to align with the demonstration scenarios of each project use case, and were subsequently demonstrated and validated as part of the evaluation of the second integrated version of the PHOENI2X framework.

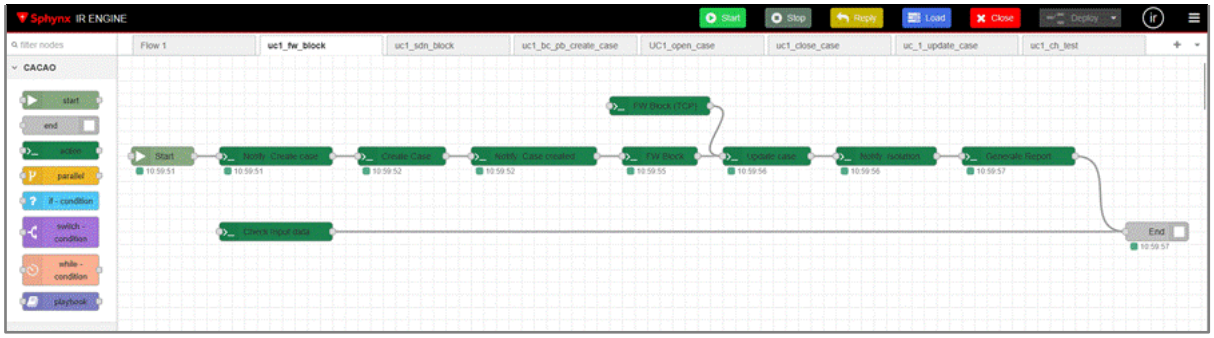

Figure 14: Business Continuity (BC) Playbook: SmartMeters Disconnection.

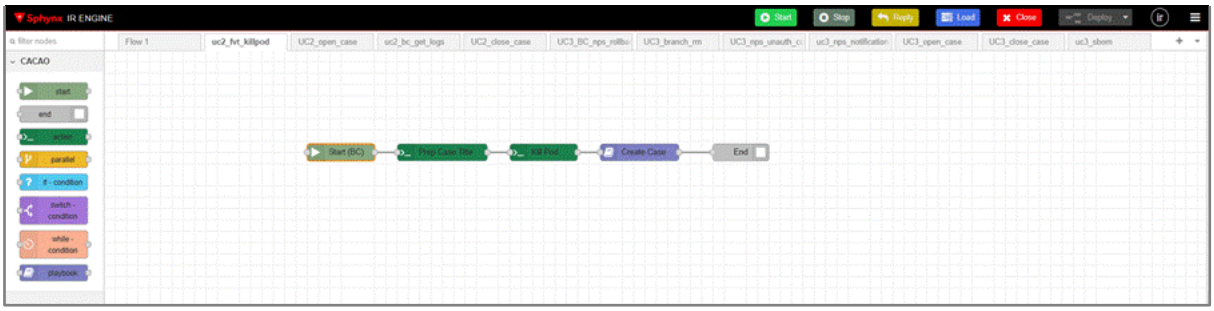

Figure 15: IR playbook – DoS attack. The playbook ends by calling the security case creation playbook.

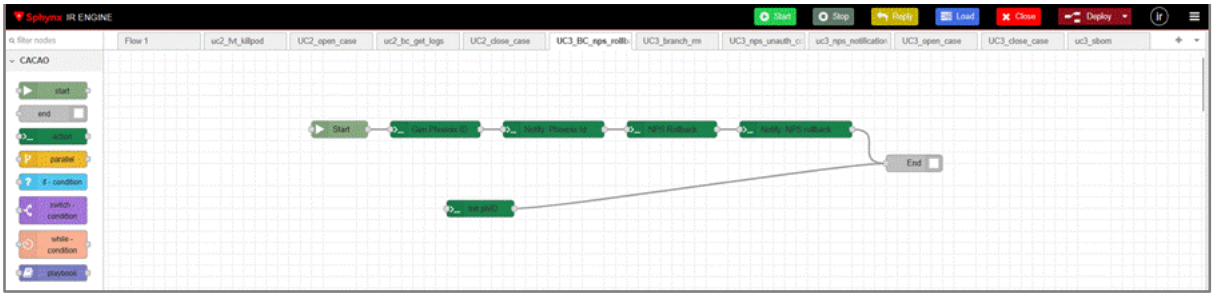

Figure 16: BC Playbook - Rollback to version without the vulnerable library.

3.4 Lessons Learned and Recommendations for Future Work

One key takeaway was the importance of modular playbook design. Creating adaptable and reusable playbooks allowed for effective customization across the Energy, Transport, and Health use cases. However, maintaining consistency and ensuring compatibility across different operational contexts required considerable effort and coordination among partners.

4 Cyber Range & Serious Games for Resilience Assessment & Training

4.1 The Resilience Cyber Range

4.1.1 Overview

The Resilience Cyber Range (RCR), developed under WP4/Task 4.3, is a core preparedness enabler in PHOENI2X, providing realistic, hands-on training in incident response, business continuity, and cybersecurity awareness.

The RCR supports Operators of Essential Services (OES) in two main ways:

1. By enabling the assessment of RPs in simulated environments that replicate real systems, helping identify and improve potential gaps.

2. By offering practical training for OES staff on business continuity, recovery, and incident response procedures defined in the RPs.

It also supports scenario-based playbook execution for "what-if" analysis and integrates trainee assessment mechanisms to evaluate preparedness. Mapping the RPs to the RCR requires configuring emulated environments accordingly.

The role of RCR, in both the high-level and detailed PHOENI2X architecture, is illustrated in Figures 17 and 18. The following sections detail the design and implementation of the RCR.

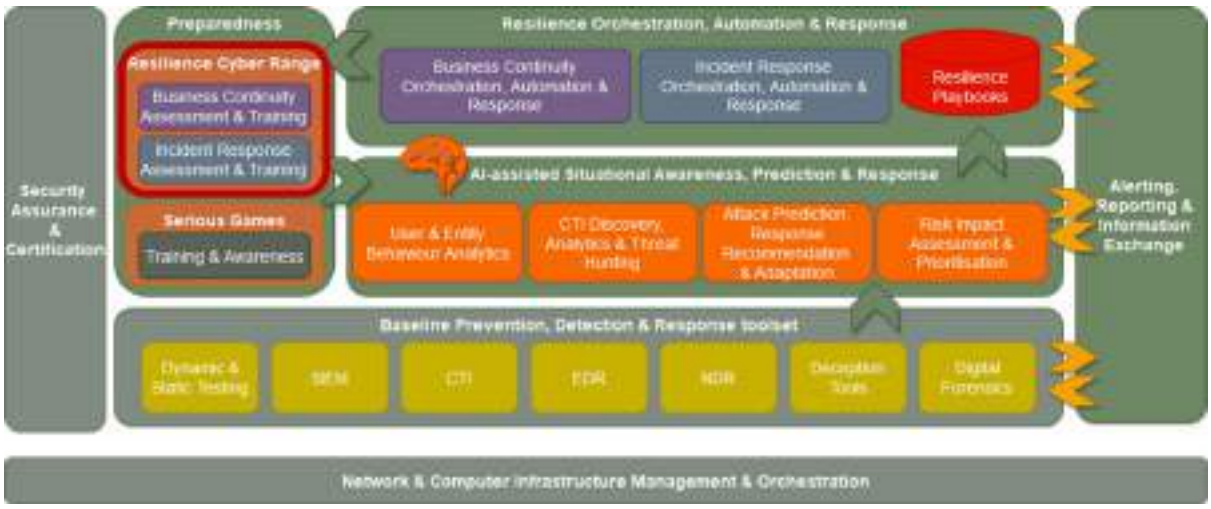

Figure 17: The RCR within the high level PHOENI2X architecture [4].

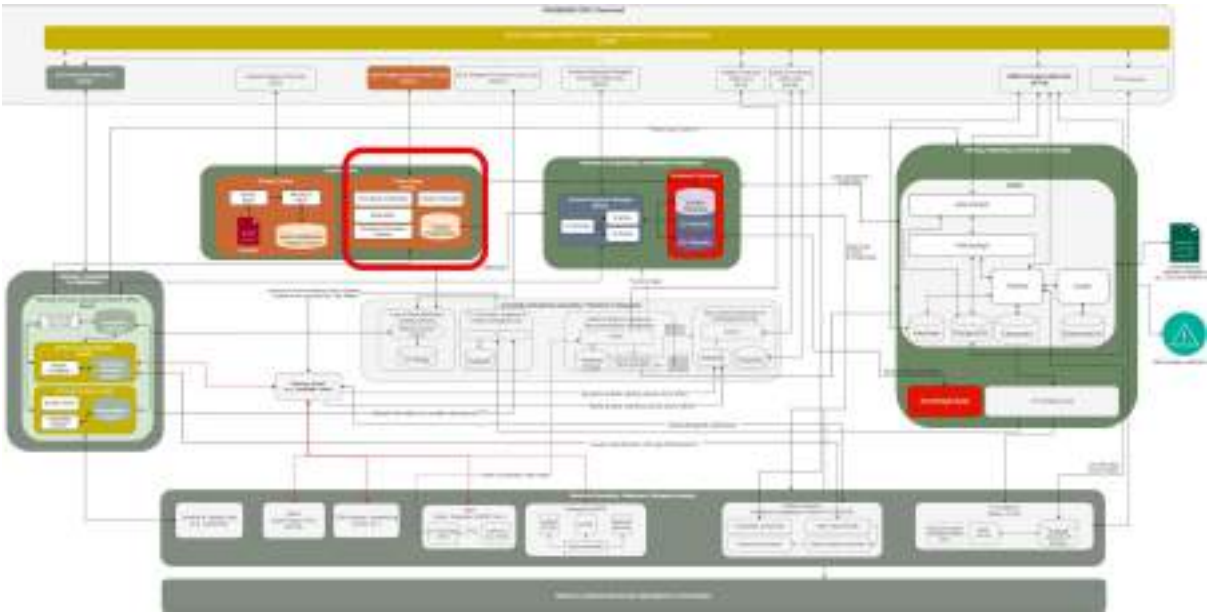

Figure 18: The RCR within the detailed PHOENI2X architecture [4]

The subsections that follow will provide more details on the current design & associated implementation of the RCR.

4.1.2 Design Details

The Resilience Cyber Range (RCR) is built on the SPHYNX Cyber Range platform, which enables the deployment of realistic, virtual training environments. It supports key user roles such as trainees, who interact with emulated scenarios, and trainers, who manage sessions and monitor progress. Training programmes and their scenarios are assigned by the training manager, who oversees the trainees' and trainers' training progress on the organisation. The platform's content, the training programmes, scenarios & emulation sandboxes are created interactively through a user interface, by the content creator.

The platform allows for the rapid setup of custom scenarios tailored to the needs of the organisation. Its modular design ensures flexibility in scenario creation and complexity adjustment.

For more technical details on the architecture, user roles, and scenario design features, please refer to Section 4.1.2 of Deliverable D4.1.

4.1.3 Implementation Details

On a high-level overview of its architecture, the RCR platform consists of four main components: the REST API, the Emulation Engine (EE), the Progression Engine (PE) and the Communication Module (CM). For more details and the relevant screenshots of each component please refer to D4.1 and in section 4.1.3.

As the project concludes, the full set of capabilities envisioned for the Cyber Range (CR) has been successfully implemented. These include support for transitioning from playbooks to training programmes, dedicated business continuity training scenarios, and the generation of tailored training programmes aligned with the PHOENI2X use case environments. All planned functionalities have been developed, integrated, and demonstrated during the second cycle of the project, along with the execution of the corresponding training programmes.

The figures below present some indicative updates to the RCR user interface. In particular, Figures 19 and 20 provide a visual overview of how a training scenario is created within the CR environment and how the progression of the training can be structured and customized. These updates aim to enhance the user experience and improve the overall design process for training scenarios.

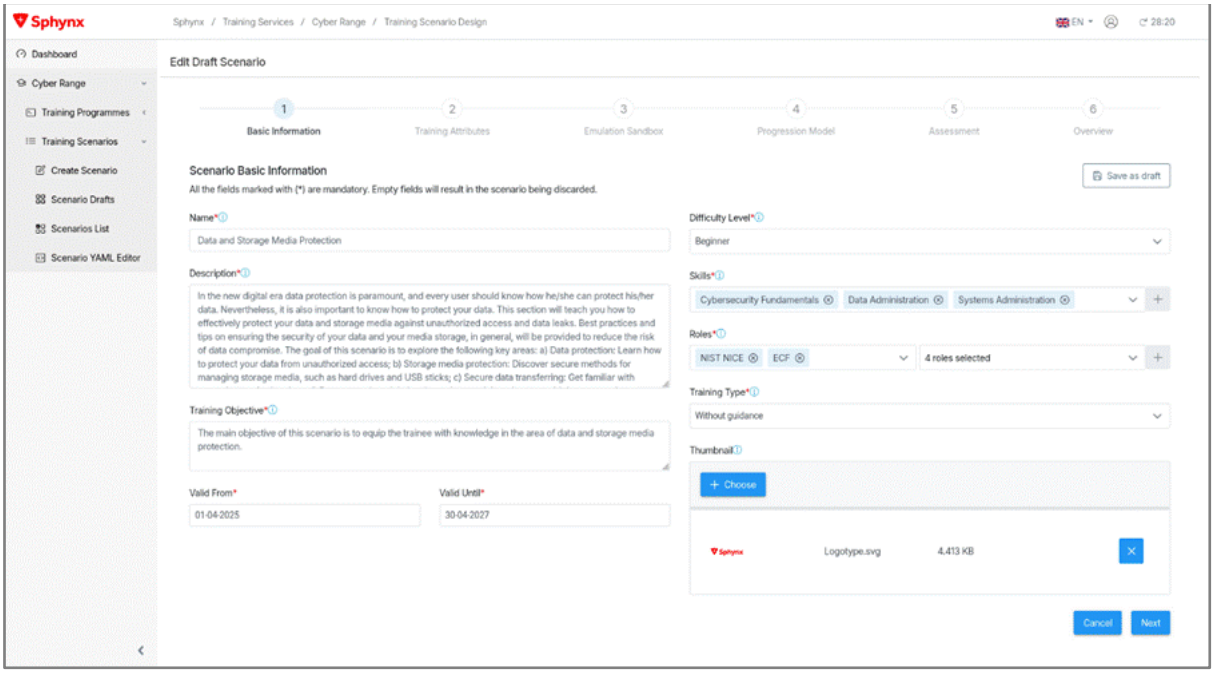

Figure 19: Creating a new training scenario.

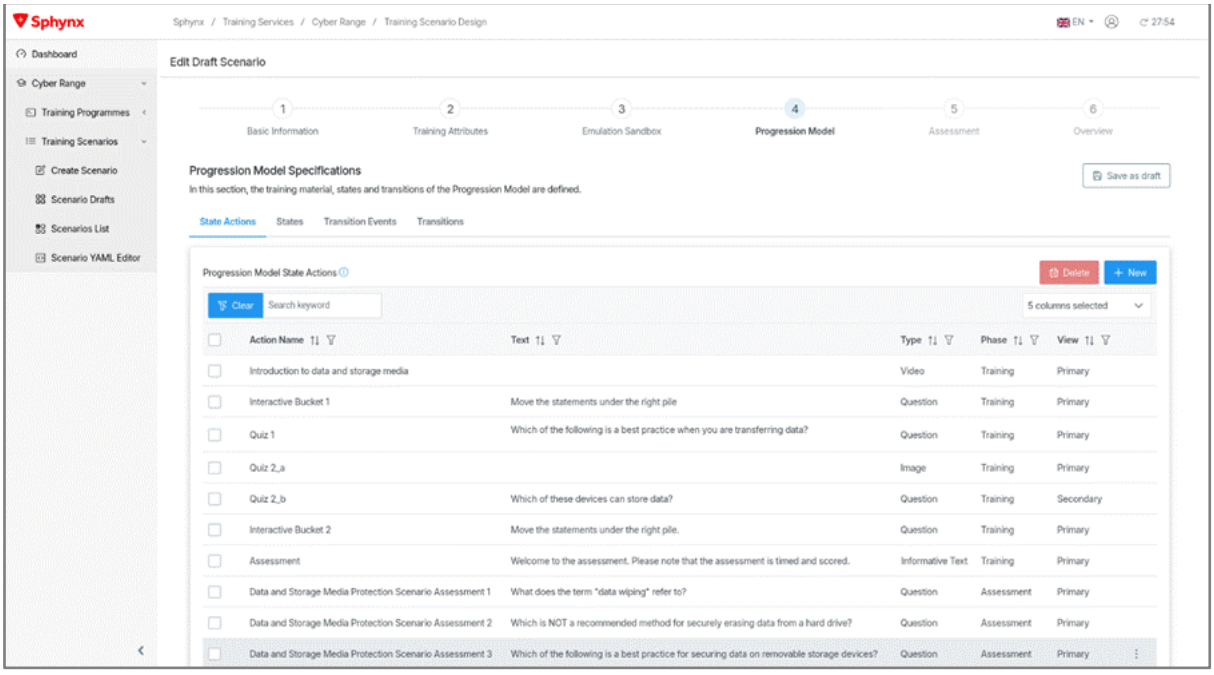

Figure 20: Designing the training scenario Progression.

Assuming multiple training scenarios are available in the system, Figure 21 showcases the updated interface that allows users to view, filter, and select from the available training programmes on the SPHYNX Cyber Range. This improved UI not only streamlines the selection process but also offers a comprehensive overview of each training programme. Key information displayed includes the programme title, difficulty level, duration, validity period, and training completion status. Additionally, users can view detailed descriptions of the training content, monitor the progress of individual training scenarios (e.g., “In Progress”, “Completed”), and read or leave user feedback through the built-in rating and comment system. A dedicated section also outlines the programme's learning objectives and assessment criteria, enabling users to make well-informed decisions based on both content and performance expectations.

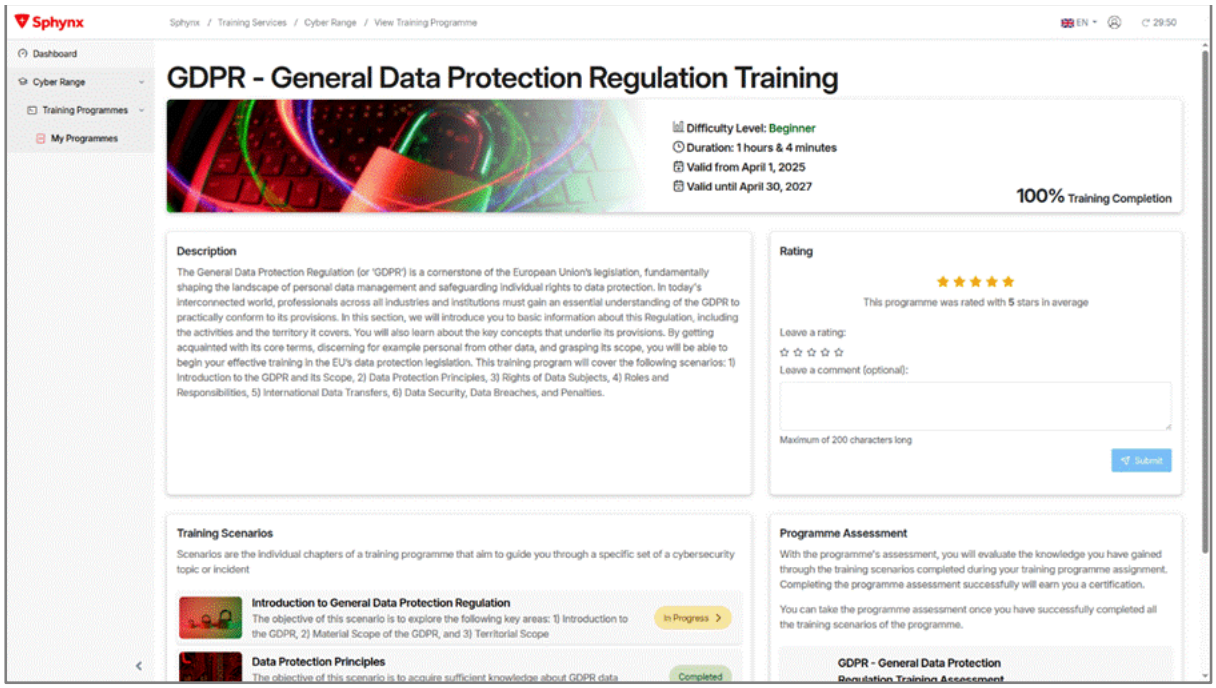

Figure 21: A training programme on the Sphynx Cyber Range.

Figure 22 shows the “Password Protection” training scenario on the SPHYNX Cyber Range. Aimed at beginner users, this theoretical, self-guided module focuses on defensive cybersecurity skills. It covers key topics such as password creation, secure storage, and password management tools. The interface displays the scenario's status, number of assessment attempts, and completion date. Supporting material from ENISA are included, and the scenario targets skills in cybersecurity awareness and management, relevant to roles like IT Project Manager and Information Systems Security Manager.

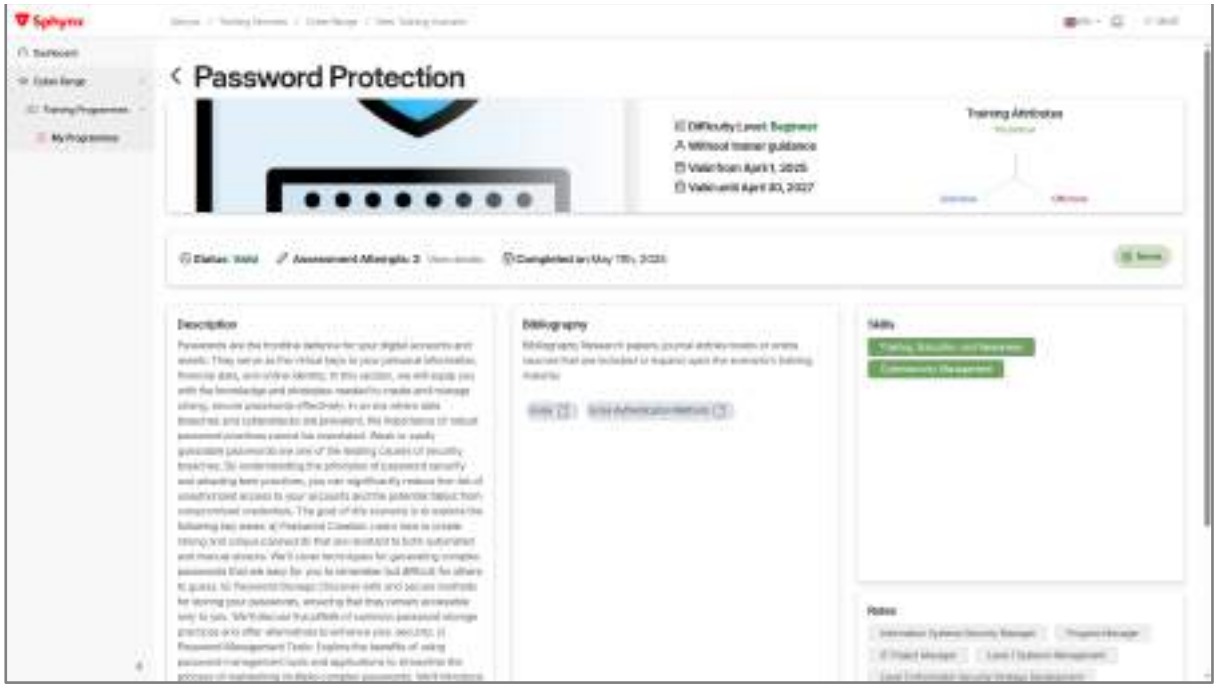

Figure 22: A training scenario on the Sphynx Cyber Range.

Additionally, screenshots of the CR interfaces during scenario execution are presented in Figures 23 and 24. Figure 23 shows the Trainee view, the interface used by the trainee to interact with the programme, submit inputs, and monitor key indicators such as elapsed time and score. Figure 24 shows the Trainer view, which displays the progression graph of the training scenario. This graph indicates the trainee's current state and provides an overview of the entire training flow.

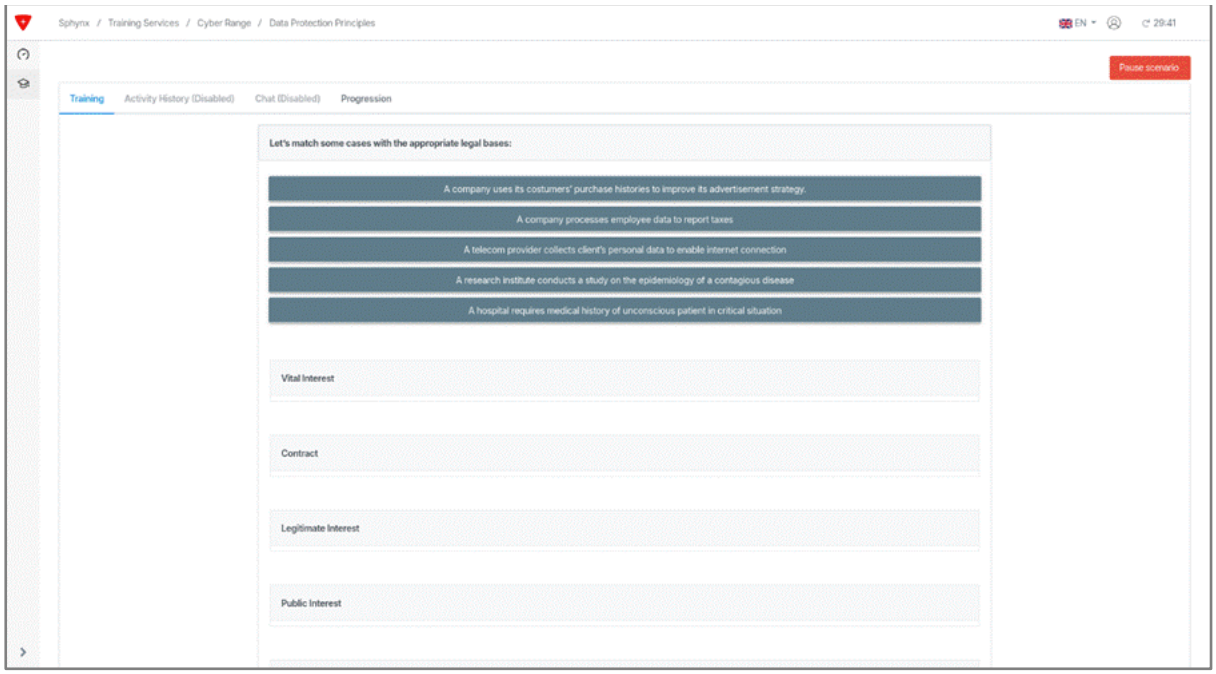

Figure 23: Execution of a training scenario.

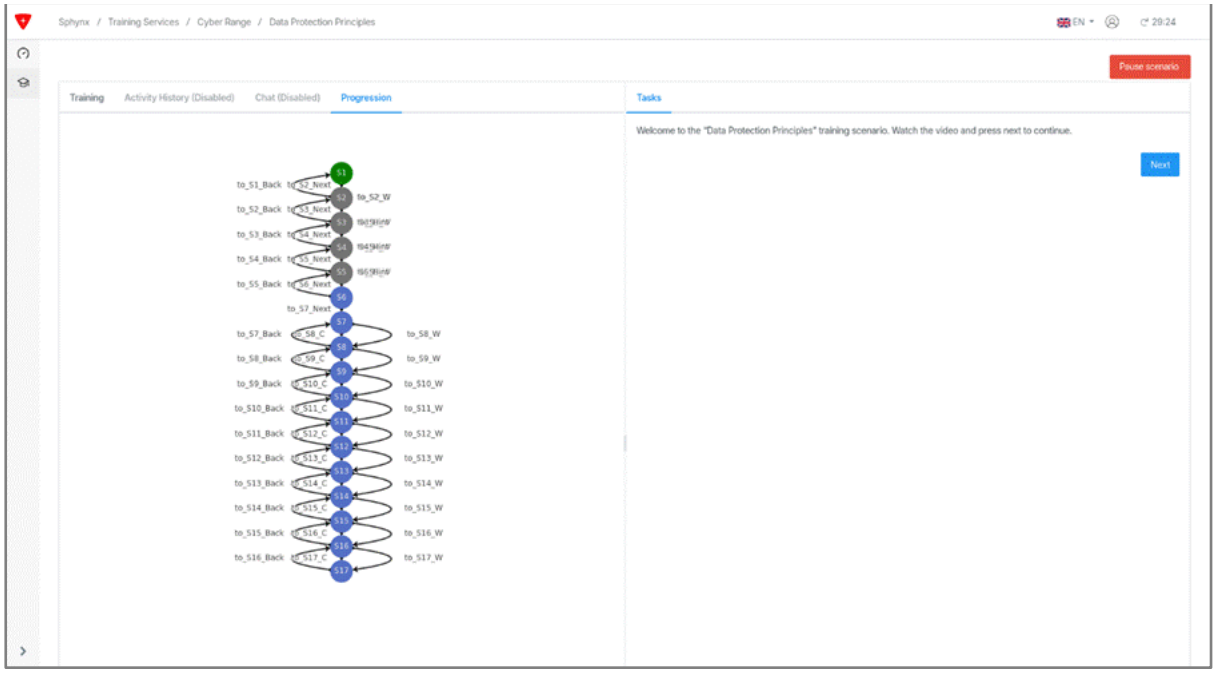

Figure 24: The FSM-based progression engine of the CR.

Finally, Figure 25 presents the end screen displayed to the trainee upon completion of the training programme. This screen provides a summary of the trainee's performance, including the final score, number of attempts, and overall result (e.g., 'Failed'). Additional performance metrics are shown, such as Accuracy and Answered percentages, accompanied by visual indicators. It also breaks down the trainee's answers per question, highlighting which questions were answered correctly, incorrectly, or left unanswered. This detailed feedback helps the trainees identify areas for improvement and track their progress over multiple attempts.

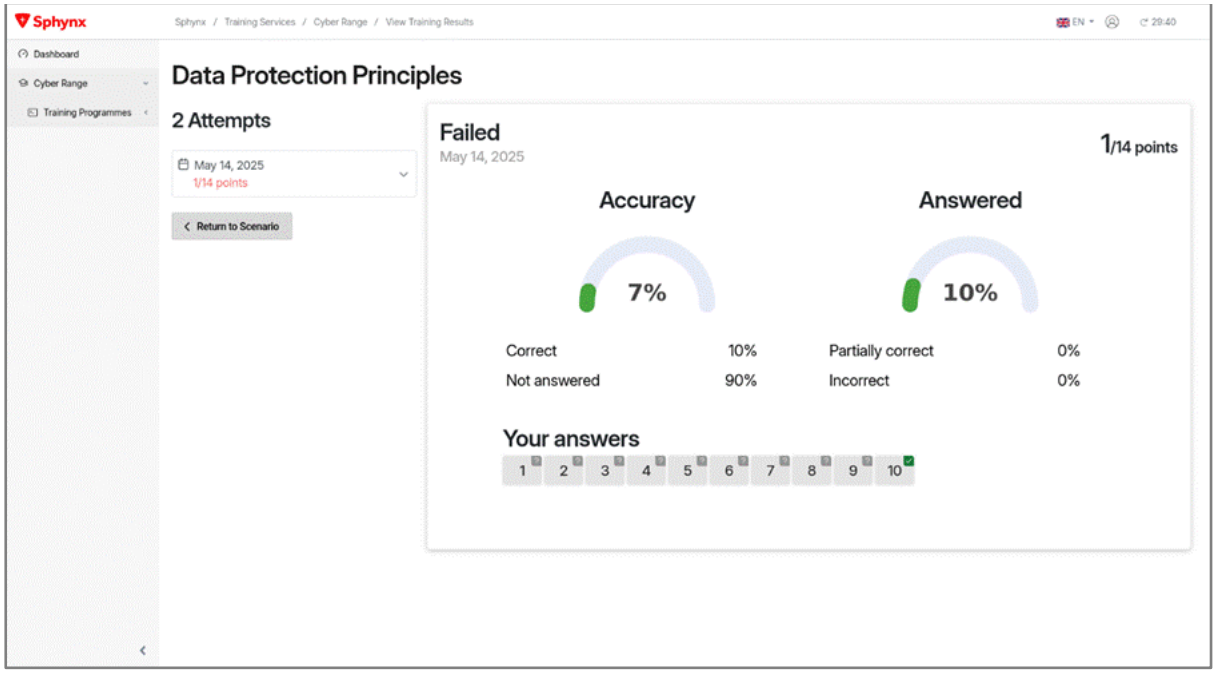

Figure 25: Scenario attempts & scoring.

Figure 26 displays the user dashboard on the SPHYNX Cyber Range platform. It provides an overview of the trainee's assignments, progress, skills, and attendance. Key metrics include the number of assigned and completed training programmes, scenario count, progress percentage, success rate, and average score. The dashboard also visualises acquired skills and tracks time spent on training, helping users monitor their learning journey effectively.

Figure 26: User Dashboard.

Figure 27 presents the Training Manager Dashboard. It provides an overview of available training programmes and scenarios, along with detailed assignment statistics for the selected programme. A pie chart visualises user progress, categorising assignees as “Not started yet”, “In progress”, “Successfully completed”, or “Assessment failed”, helping managers monitor training engagement and performance.

Figure 27: Training Manager Dashboard.

Figure 28 displays the User Profile page. It shows user information, including name, role, email, and membership date. The page highlights acquired skills through interactive radar charts covering theoretical, defensive, and offensive cybersecurity competencies. It also lists the user's current training assignments, offering a personalised overview of progress and skill development.

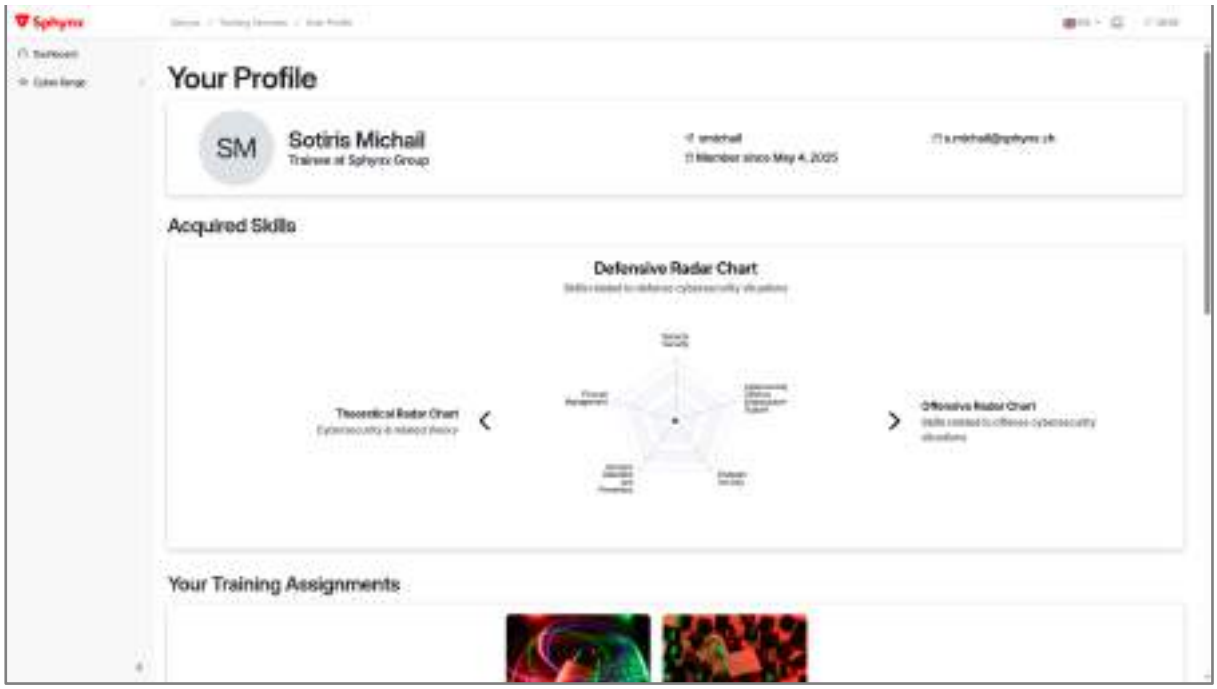

Figure 28: User Profile.

As the baseline capabilities of the Cyber Range (CR) were implemented during the previous cycle of the project, this second cycle focuses on extending its functionality to better support PHOENI2X and its specific requirements. These enhancements include features such as trainer view support, more detailed reporting on training progress and status, support for transitioning from playbooks to training programmes, business continuity training capabilities, and the generation of training programmes tailored to PHOENI2X use-case environments.

4.1.4 Final Status

During the second cycle of the project, the Resilience Cyber Range (RCR) environment was finalized and deployed in a cloud-based infrastructure. This transition to the cloud was motivated by the need for improved scalability, ease of maintenance, and broader accessibility for training participants across different locations.

Access to the Cyber Range (CR) will be granted upon user request. This controlled access approach ensures that only authorised participants can utilise the platform, supporting proper scheduling, user management, and resource allocation. It also allows the administrators to monitor usage, maintain system integrity, and provide technical support as needed, thereby enhancing the overall efficiency and security of the training environment.

Building on the baseline functionalities developed during the first cycle, the RCR has been significantly enhanced to support the execution of a wide range of training programmes. It is now capable of emulating complex environments and scenarios tailored to both general and sector-specific requirements.

For the purposes of PHOENI2X, a comprehensive and diverse set of training modules has been developed to address a wide spectrum of user needs and technical competencies. These modules cover general cybersecurity awareness topics (such as Cyber Hygiene and GDPR compliance) aimed at non-technical users and organisational staff. In addition, more advanced technical programmes have been created, including training in areas like Incident Response and Supply Chain Security. Beyond general and technical content, specialised training scenarios have been designed to meet the specific requirements of key critical infrastructure sectors, namely Energy, Transport, and Health. These sector-specific trainings emulate realistic environments and incidents, providing trainees with hands-on experience tailored to the operational challenges and threat landscapes of each domain.

Where applicable, the Resilience Cyber Range (RCR) has been integrated with incident response playbooks developed within the PHOENI2X project. This integration enables scenario-based training sessions that simulate realistic incidents and reinforce appropriate incident response workflows.

Available Scenarios within the Cyber Hygiene Programme:

- Phishing
- Passwords
- Data & Storage Media Protection
- Remote Working & Leisure
- Secure Software Configuration

Available Scenarios within the GDPR Programme:

- Data Protection Principles
- Rights of Data Subjects
- GDPR Roles & Responsibilities
- International Data Transfers
- Data Security, Data Breaches, and Penalties under the GDPR
- Training contents under integration

Additional Training Programs:

- Generic Incident Response Training
- Sector-Specific IR Training (Energy): Detection of DLMS/COSEM Attacks
- Sector-Specific Training (Energy/Transport): NIS2 Directive
- Sector-Specific Training (Healthcare): NIS Directive
- Sector-Specific Training (Energy): Supply Chain Security and Secure Coding

4.1.5 Lessons Learned and Recommendations for Future Work

The development and deployment of the Resilience Cyber Range (RCR) provided several key lessons.

One of the main lessons-learned was the importance of flexibility and modularity in scenario design. Creating a training environment capable of supporting diverse sector-specific needs (e.g., Energy, Transport, and Health) required modular templates and configurable infrastructure components. Future versions of the RCR could benefit from further automation in scenario deployment and dynamic environment generation to streamline the preparation of training exercises.

Another important takeaway was the value of moving the RCR to a cloud-based setup, which greatly improved accessibility and scalability. However, challenges around ensuring consistent performance across different user locations and isolated environments for each training session, remain areas for improvement. Future work could focus on optimizing cloud resource management and incorporating orchestration tools for smoother operation.

A challenge that emerged was the integration of assessment mechanisms to track trainee performance during exercises. While some basic functionality was implemented, future development could include more advanced analytics, automated scoring, and adaptive training paths based on trainee behavior and response times.

Additionally, establishing a tight integration between incident response playbooks and the RCR, proved highly valuable, enabling hands-on exercises that reflect real-world workflows. This connection could be further enhanced by adding feedback loops, where the outcome of a training session informs about updates to the playbook or the RCR scenario itself.

Lastly, collaboration with domain experts and consortium partners was crucial for designing realistic and impactful scenarios. Future work could benefit from developing a library of reusable training modules and launching a community-driven approach to continuously improve scenario quality and relevance.

4.2 Serious Games

The concept of serious games within the PHOENI2X framework is a module that provides an interactive, context-specific learning experience. The initiative encompasses a series of games[2], each with a distinct focus: from the physical tabletop game HATCH, designed to foster understanding and mitigation strategies against social engineering attacks among domain experts, to digital games like PROTECT and AWARENESS QUIZ, which broaden the scope to include a wider audience. The concept is to apply that learning approach in dynamic and realistic scenarios, augmented by machine learning technologies to personalize and enhance the learning experience. This approach not only aims to increase cybersecurity awareness but also to refine the participants' ability to identify and defend against real-world security threats.

4.2.1 Overview

PHOENI2X project approach contains serious games that have the following goals:

- Context-specific Cybersecurity Awareness Training.
- Context-specific Social Engineering Threat elicitation and assessment.
- Simple, engaging and target training methods.
- Training content based on real world security incidents.
- Engaging graphical representation.

The gamification approach consists of one physical tabletop game that engages domain experts in an adventure to design social engineering attacks that work in their specific context. The proposed attacks are rated by the other players, who are also domain experts. The game is moderated by multiple cybersecurity experts to ensure a realistic output of threats. The tabletop game is played with a sample of stakeholders from the relevant context. By experience, playing with more players does not reveal a significant number of new threats after a certain threshold of players is reached. In addition, the game motivates the players to design realistic attacks to win the game. This triggers their motivation for learning cybersecurity.

After the threats have been elicited with the tabletop game HATCH, we use the elicited threats in a digital card game called PROTECT in which the target audience is given attack and defense cards. To win the game, they have to choose the correct defense to a social engineering attack derived from the elicited threats. Finally, the digital game AWARENESS QUIZ presents real attacks related to the elicited threats urging players to respond to this realistic scenarios.

The digital games have been augmented with ML technologies to:

- Adapt the difficulty of the game.
- Re-introduce attack/defense pairs that the player got wrong.
- Choose a set of questions in the quiz related to the threats that the player got wrong before.
- Support the player in replying to cybersecurity risk attitude test.
- Evaluate risk attitude test and recommend training content.
- Search for context-specific attacks on the web to create PROTECT card decks and quiz questions for AWARENESS QUIZ.

[2] https://www.social-engineering.academy/en/offers.html

4.2.2 Design Details

The design details of these serious games reveals a comprehensive structure tailored to the intricacies of social engineering threats. The HATCH game, for instance, provides an immersive experience that helps players recognize and counteract social engineering tactics in a controlled, game-based environment. This hands-on approach is supplemented by the PROTECT and AWARENESS QUIZ digital games, which extend the learning experience to include defense mechanisms against digital threats and to reinforce the importance of cybersecurity awareness through interactive quizzes. These games are meticulously crafted to ensure the educational value, the engagement and motivation for the players, leveraging real-world scenarios and threats to impart crucial cybersecurity knowledge and skills.

4.2.3 Tabletop Game: HATCH

In addition to IT infrastructure, employees are also a target by cybercriminals. Through social engineering attacks such as phishing emails, cybercriminals try to exploit the human characteristics of the participants in order to tempt them to perform harmful actions.

The serious game HATCH (Figure 29 & Figure 30) provides an entertaining, interactive group training that teaches employees to identify and successfully defend against social engineering attacks.

Figure 29: Main game board of game HATCH.

With the help of serious games employees can be engaged in security activities in an enjoyable and sustainable way, to increase the awareness of (and defensive behaviour against) social engineering threats. The trainings are conducted exclusively by cybersecurity experts.

Figure 30: HATCH Gender Inclusive Persona Cards.

- In the first cycle the focus was on the adaptation of the scenario. In the second cycle, the game board and the persona cards were adapted to be gender inclusive by offering a female and a male persona to the players. Further work was the investigation of the usage of generative AI for the newly created railways / transportation scenario. Complies with ISO 27001 Control A.7.2.2.
- Company-specific adaptation (e.g. industry sector) of the game scenario is possible.
- Identification and assessment of relevant social engineering threats.

4.2.4 PROTECT, Online Game

The serious online game PROTECT (Figure 31 & Figure 32) enables a motivating and entertaining training that actively sensitizes employees against social engineering threats.

Figure 31: Main game board of game PROTECT.

Training should not be boring. With the help of serious games, employees can be engaged in security activities in an entertaining and sustainable way.

Figure 32: PROTECT Sample Cards.

Key features of PROTECT include the following.

Relevant learning content

- Customization of content to your security policies possible
- Complies with ISO 27001 Control A.7.2.2
- Multiple languages

Online provision

- Time and location independent training
- Scalable cloud solution
- Playable on PC and mobile devices

Performance measurement

- Data protection compliant evaluation
- High score leader board
- Certificates for participants

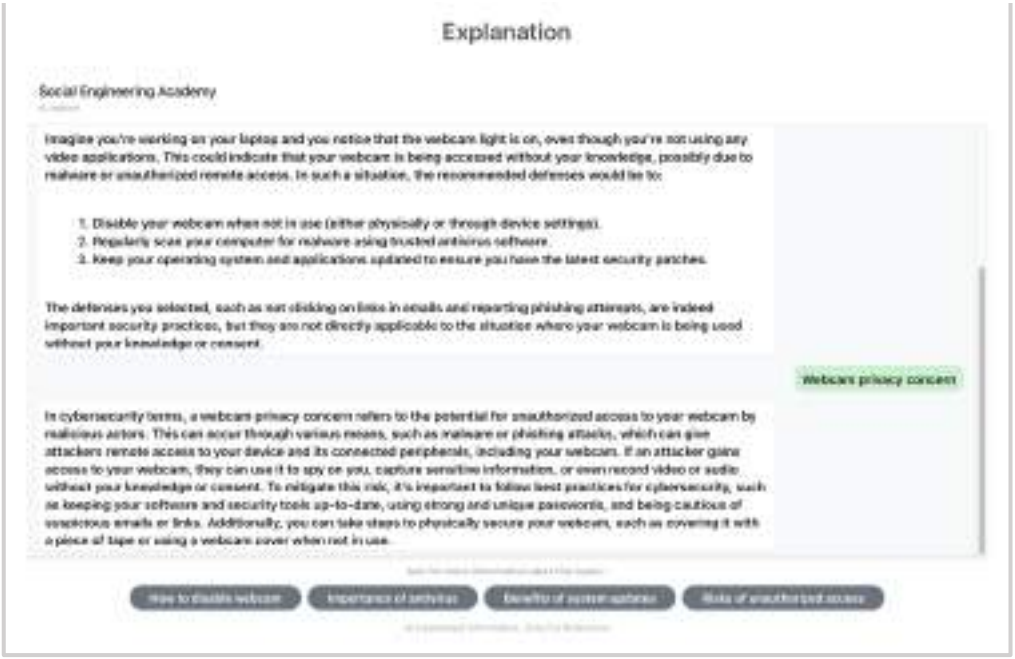

Figure 33: PROTECT LLMs Explanation Sample.

Large Language Models Assistance in the Game

In our serious game, players engage in a card-matching mechanic where they must correctly pair attack cards with the most appropriate defense cards. This matching task is central to the game's learning objective.

When a player selects an incorrect defense for a given attack, a Large Language Model (LLM) is activated to provide immediate, contextual feedback. The system takes both the current attack card and the wrongly selected defense card and submits them to the LLM. The model then explains why the selection is incorrect, helping the player understand the reasoning and deepen their learning.

Key Features:

- LLM Location: The LLM is hosted within the European Union (France).
- Privacy Compliance: The provider is fully GDPR-compliant, ensuring user data is handled responsibly and securely.
- Extended Learning Through Guided Dialogue:

  To prevent unrestricted use of the LLM and maintain focus, the system proposes up to five follow-up questions related to the topic of the mistake. These are context-aware and crafted to help the user explore the subject further. Players can click on these questions to continue a guided conversation with the model, promoting deeper comprehension of misunderstood concepts.

This integration of AI assistance transforms incorrect answers into valuable learning opportunities, enhancing both engagement and educational outcomes.

4.2.5 AWARENESS QUIZ, Online Game

The serious game CYBERSECURITY AWARENESS QUIZ (Figure 34) represents an online quiz that sensitizes the participants for social engineering threats and their possible effects.

Figure 34: Main game board of game AWARENESS QUIZ.

The interactive and hands-on training with the online quiz CYBERSECURITY AWARENESS QUIZ imparts awareness with respect to the potential impact of social engineering attacks. The quiz content is based on real attacks, for which the sources are also provided. This creates a direct reference to the everyday work of employees and highlights the relevance of the content.

Key features of AWARENESS include the following:

Relevant learning content

- Based on real attacks
- Constant updating
- Customization of content to suit your organization
- Complies with ISO 27001 Control A.7.2.2
- Multiple languages

Motivation

- Predefined quizzes with relevant topics
- Composition of individual quizzes
- Various multiplayer modes
- High score leader board

Online provision

- Time and location independent training
- Scalable cloud solution
- Playable on PC and mobile devices

SEA contributed two serious games for educating people about the dangers of social engineering attacks specific to the scenarios of PHOENI2X. Further, the tool receives exercise configuration data from the PHOENI2X platform to:

- Configurations for the run time of a game
- Data about attacks to be integrated in the games
- Difficulty level of a game
- Definition of further game settings

Technically, we concise the digital games in a platform called Gamification Tool. In the Gamification Tool, two main games are currently documented by SEA (PROTECT and AWARENESS QUIZ) and are included in the project. The gameplay of the game PROTECT was further improved. This web-based game will be the first game that will be integrated into the platform. The initial steps towards its visualization are also documented in this period. Additionally, SEA documented the tabletop card game HATCH and improved the design as well as the content of the cards to align with the PHOENI2X use cases. Regarding HATCH, SEA also created specific persona cards and a game plan for the energy scenario. Training sessions with this scenario have been provided by SEA during the last project meeting (January 2024) in Athens.

Moreover, an initial analysis is performed on the tools' communications and data flow under the umbrella of the platform architecture activities. The results of the platform's architecture form the baseline of the activities and define the necessary communication mechanisms and interfaces for information exchange between the Training and Visualization Tools and the rest of the PHOENI2X components. Further analysis of the data format, the communication security needs (e.g. the use of Transport Layer Security, TLS) and type of communications, such as synchronous or asynchronous, has been performed for the needs of this deliverable.

Currently, data communication to and from the ML engine are being designed.

4.2.6 Implementation Details

Figure 38 shows the Gamification Tool architecture high-level view. Two serious games are currently offered as part of the Gamification Tool component.

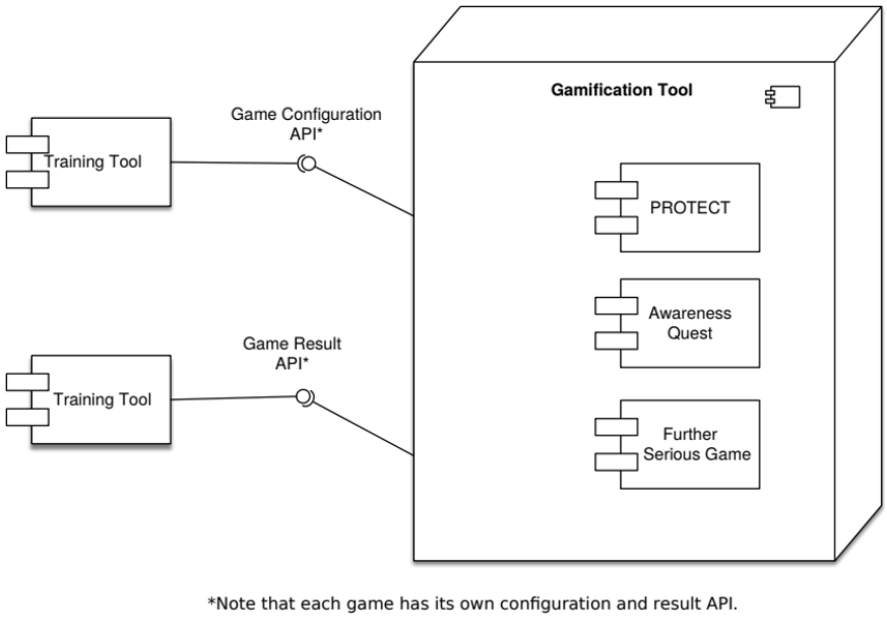

Figure 35: Gamification Tool Architecture and Message Flow.

The serious games are started by a user via their own GUI. They are developed using Angular 6[3] and use REST for all software interfaces. The players can play them on their PC browsers or on a mobile device. The games support automatic scaling to the type of screen size they are played on. The games each have their own interface, because the data needed to configure them highly depends on the game itself and is almost not generalizable.

4.2.7 Gamification Tool Interfaces

Conceptually, there are two APIs provided/required: one for configuring games e.g. run-time, difficulty, userid etc. by the Training Tool and one for transmitting the results and the userid to the Training Tool.

Table 1: Gamification Tool – Protect Interfaces.

Interface/Operation Name	Input Data	Output Data	Description
GameSetting.Protect.playerID	Integer		Unique ID of player
GameSetting.Protect.playerName	String		Name of the player
GameSetting.Protect.gameTime	Integer		Time the game is running
GameConfiguration.Protect.difficulty	Integer		Difficulty level
GameConfiguration.Protect.jokers	Integer		Number of jokers in the game
GameConfiguration.Protect.attacks	String		New attack and defense description
GameResult.Protect.score		Integer	Points scored in the game
GameResult.Protect.highscore		String	Entire high score list of all players
GameResult.Protect.playerID		Integer	Unique ID of player
GameResult.Protect.playerName		String	Name of the player

[3] https://angular.io/

Table 2: Gamification Tool – Awareness Quiz Interfaces.

Interface/Operation Name	Input Data	Output Data	Description
GameSetting.AwarenessQuiz.playerID	Integer		Unique ID of player
GameSetting.AwarenessQuiz.playerName	String		Name of the player
GameSetting.AwarenessQuiz.gameTime	Integer		Time the player can answer questions
GameConfiguration.AwarenessQuiz.challenge	Integer		Complexity level of the questions
GameConfiguration.AwarenessQuiz.Questions	String		Add new questions and answers to the game
GameResult.AwarenessQuiz.Correct		Integer	The number of correct answers a player provided
GameResult.AwarenessQuiz.playerID		Integer	Unique ID of player
GameResult.AwarenessQuiz.playerName		String	Name of the player

We specified the interfaces of the serious games PROTECT and AWARENESS QUIZ in Table 2 and Table 3, respectively. We use REST services for data exchange, communicating in a JSON data format.

The games PROTECT and AWARENSS QUIZ are programmed in Vue.js and are developed and communicated over REST web services using messages in the JSON format. All configurations and player data can be set and retrieved using REST and JSON. We selected common technologies that are used widespread to ensure compatibility with numerous end user devices and common standards for data exchange in web applications.

Note that these games do not store any user information after a game is finished. When a game is finished, the game reports userid and game result back to the training tool.

We have the following assumptions regarding the PHOENI2X platform:

- Identification of relevant staff for training and Access Management is provided by the PHOENI2X platform.
- GDPR compliance for all user data (e.g. consent of players to process their data) is provided by the platform.

5 Alerting, Reporting & Information Exchange

5.1 Smart Mandatory Incident Reporting Tool (SMIR)

The Smart Mandatory Incident Reporting Tool (SMIR) included in the Alerting, Reporting & Information Exchange capabilities of PHOENI2X provides the Critical Infrastructures' incident reporting teams with a solution to support them during the mandatory incident reporting process for compliance with the applicable regulatory frameworks. As described in detail in section 5.1.1 of D4.1 [4], SMIR is an evolution of the AIRE asset developed in the project CyberSec4Europe[4]. Its architecture is the same detailed in section 5.1.2 of D4.1. In summary, it is composed of different services, where the main ones are the ***aire-workflow-enforcement,*** which is the core module responsible for the enforcement of the incident reporting workflow, and the ***aire-reports-generator***, which is in charge of the generation of the incident reports according to different templates. SMIR is integrated with the open-source Security Incident Response Platform TheHive[5] through the service ***aire-thehive-plugin***.

However, the incident reporting workflow described in section 5.1.2 of D4.1 has been extended to integrate the improvements developed in SMIR in the context of the H2020 European project SUNRISE[6] to add support for compliance with the new NIS2 Directive, which are also applicable to PHOENI2X. Current Business Process Model and Notation (BPMN) enforced by SMIR is the one shown in Figure 36. In particular, it has been extended to integrate risk-based classification, include early warning of incidents initially classified as significant, notification to service beneficiaries and new procedures for countermeasures and vulnerability management. More details about this new workflow and the BPMNs used for these additional procedures are described in SUNRISE D6.5 [5].

[4] https://cybersec4europe.eu/
[5] https://thehive-project.org/
[6] https://sunrise-europe.eu/

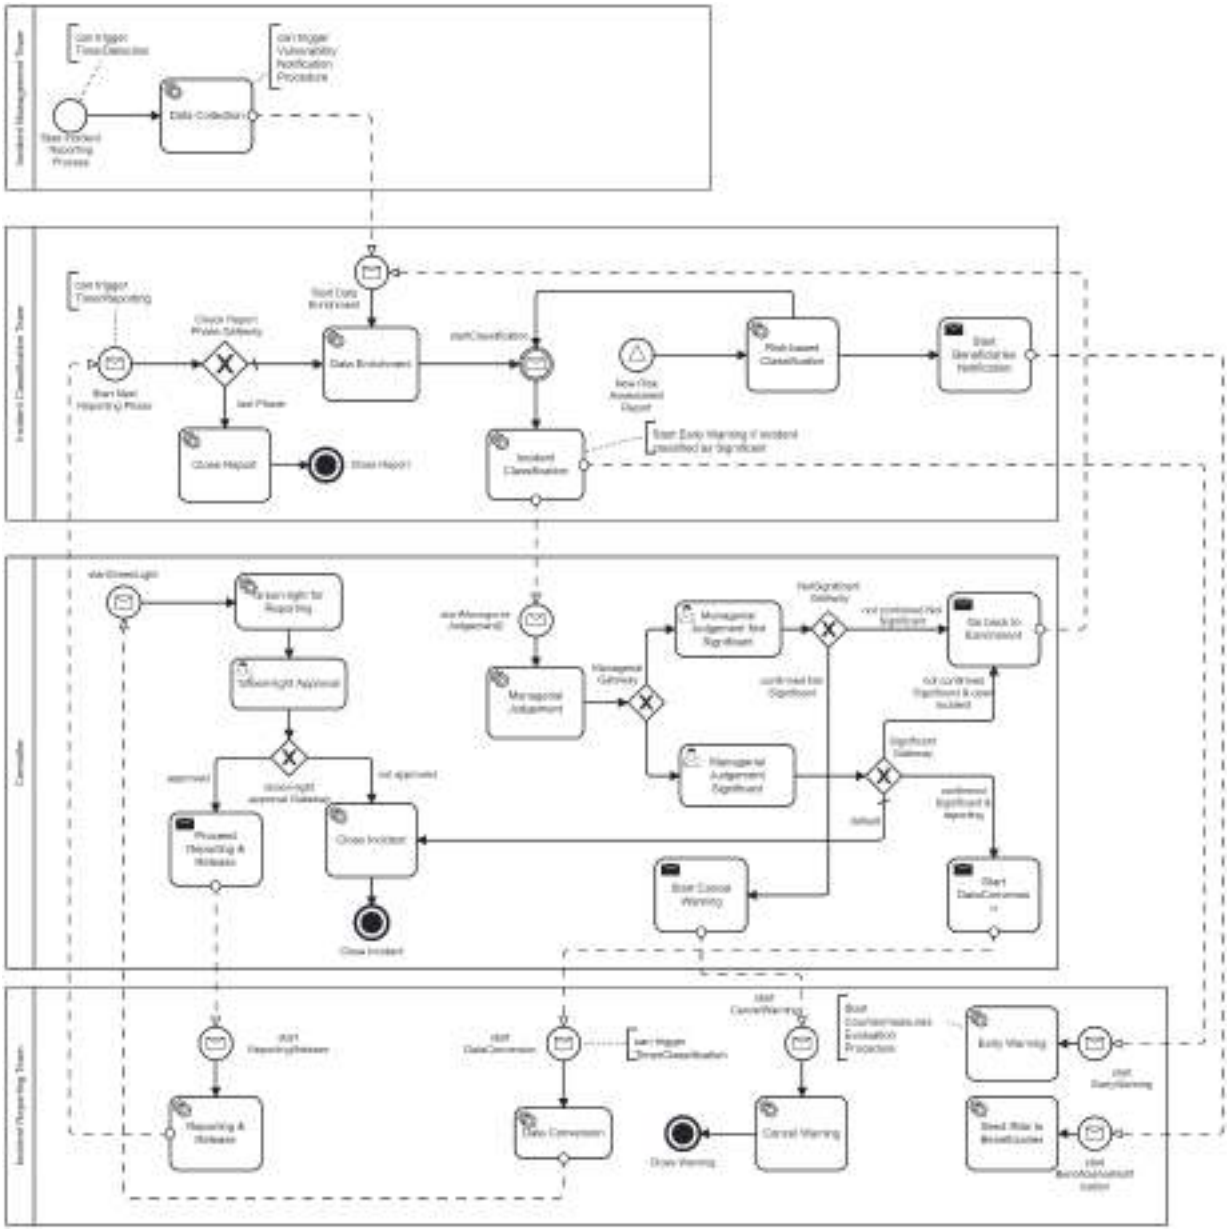

Figure 36: SMIR Incident Reporting BPMN (extracted from Figure31 of SUNRISE D6.5 [5]).

In the context of PHOENI2X project, additionally to the developments already reported in section 5.1.3 of D4.1 [4], enhancements done in the SMIR tool during this last period are presented below:

- Automation of the notification of the incidents to the Supervisory Authorities. In compliance with NIS2 directive, SMIR has been extended to support the automatic sending of an early warning when an incident is classified as significant to relevant entities and the sending of the mandatory report or reports when required depending on the regulation. SMIR dashboard includes the possibility of configuring these communication channels and optionally enabling them for automatic reporting (see Figure 37).

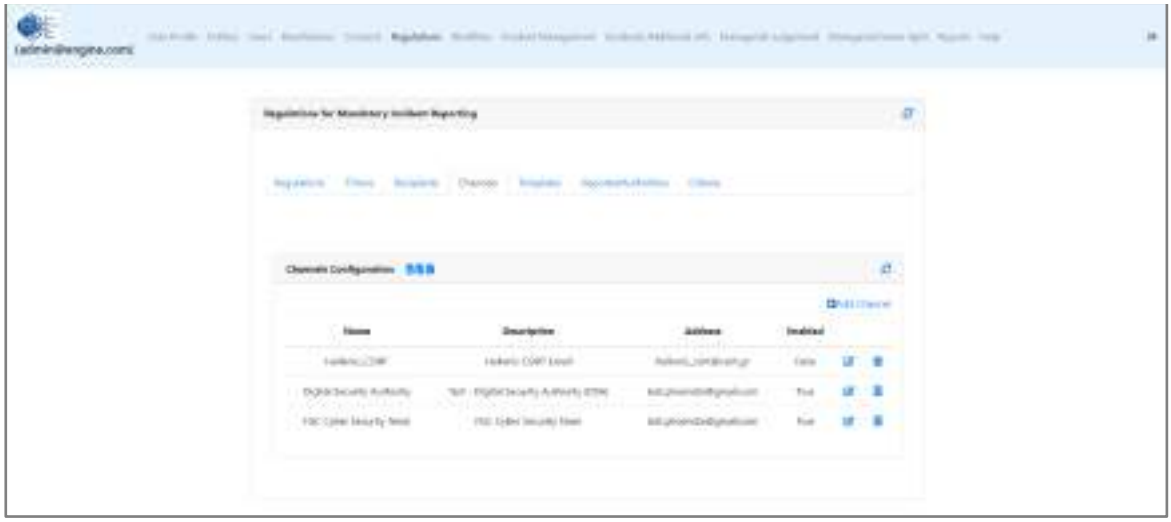

Figure 37: SMIR Dashboard – configuration of notification channels.

- Additionally to the existing endpoints described in section 5.1.3 of D4.1, a new endpoint has been added into SMIR to support the reception of requests from the ROAR to automatically send an early report for cyber security teams in the context of the Use Case 2. This new endpoint is /aire/earlyNotification/{incidentId} where incidentId must be the incident identifier returned to the ROAR when it sends a request to register a new incident in SMIR, which is indeed the id of the case registered in the open-source tool TheHive.
- Inclusion of technical information in the reports received from other components within PHOENI2X to enrich the context of the incident. This includes Indicators of Compromises that are received through the CTI component or alerts generated by the monitoring components. These alerts and IoCs can be also visualized through the graphical interface provided by TheHive and integrated or merged into the incidents as Observables. SMIR has been extended to support the inclusion of these observables in the documents generated as an annex by just configuring the template to add them. An example is shown in Figure 38 where an observable with a suspicious IP address is added to the Word document generated with the incident report for Use Case 3.

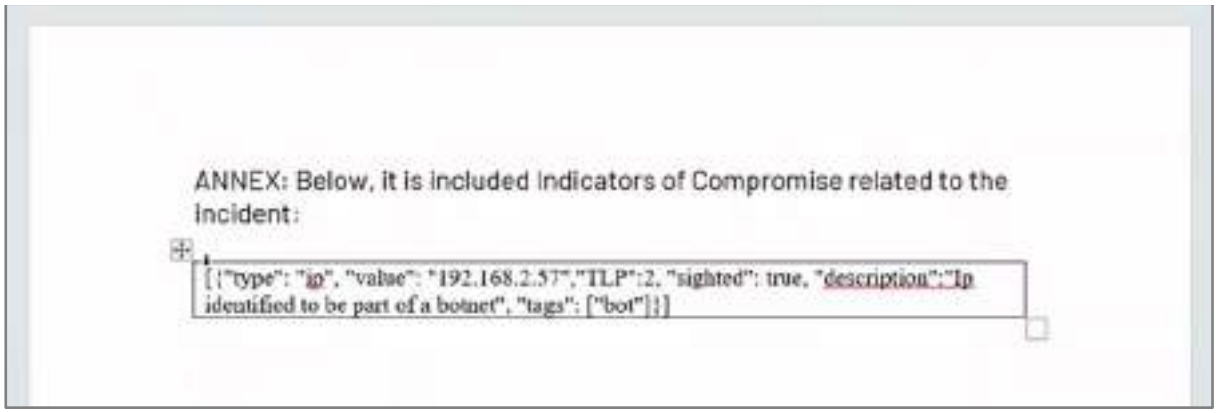

Figure 38: SMIR – Template including technical information received from CTI.

- Adaptation of SMIR to the generation of the incident report according to the template provided by the Health use case (UC3). This template is suitable for notification of incidents to the Digital Security Authority (DSA) according to the NIS2 Directive. In Figure 39, it is shown how the user can see from the SMIR dashboard that the report is available once information provided about the

incident has been converted to the DSA template. An example of the first page of this report is shown in Figure 40.

Figure 39: SMIR – Report generated for DSA according to NIS2.

TLP: AMBER

NOTIFICATION FORM OF AN INCIDENT TO THE DIGITAL SECURITY AUTHORITY

INTRODUCTORY INFORMATION - (Mandatory for all notifications)

Report by:	Kostas Lampropoulos	Telephone no:	+30123456789
Title:	Data Protection Officer	E-mail Address:	kostaskostas@dei.gr
Signature:		Report to:	

Type of notification:	Initial	Interim	Final
Notification Type:	Mandatory	Voluntary	

INFORMATION REGARDING THE INCIDENT (MANDATORY FOR ALL NOTIFICATIONS)			
Organization Name:	UPAT	Internal incident number:	UPAT_001
Essential services affected:	Health	Date and time of detection:	18/02/2025 09:30
		Date and time of submission of the report.:	21/02/2025 06:37:28
Type of incident: (Cybersecurity / Non-Cybersecurity / both)	Cyber Security Incident		
Current incident status: (Detected/probable)	Detected		
Stage of incident: (In progress/ In progress but under control/ Terminated)	In progress		
Essential services that have been affected.			
Number of hours during which the continuity of the service was affected (Indicate the number of users and the duration of the incident, in hours)	120 minutes		
Geographical area of an incident	Cross-border EU		
Were there serious injuries or loss of human life as a result of the incident?	false		

Figure 40: SMIR – DSA template for health use case (UC3)

- **Aire-reports-generator**: This is the module in charge of the generation of the required incident report files following different formats and templates. It retrieves the information about the incident stored in the database and adapts it to the different report templates.

Figure 41: Incident registered in TheHive from ROAR

- Since ROAR can register new incidents in SMIR, it also needs to receive a confirmation from SMIR when each incident is closed (e.g., to progress a playbook). As such, SMIR has been extended to support notification functionality to external endpoints.
- NIS2 directive and GDPR have been considered in PHOENI2X to cover the use case needs about reporting to their respective national authorities. For NIS2, the Word template offered by the National Cyber Security Authority of Greece (see Figure 40) has been integrated into SMIR. For GDPR, an Excel template has been created based on the information required by FGC Cybersecurity procedures and regulations (see Figure 50). The templates can be configured via the GUI (see Figure 51).

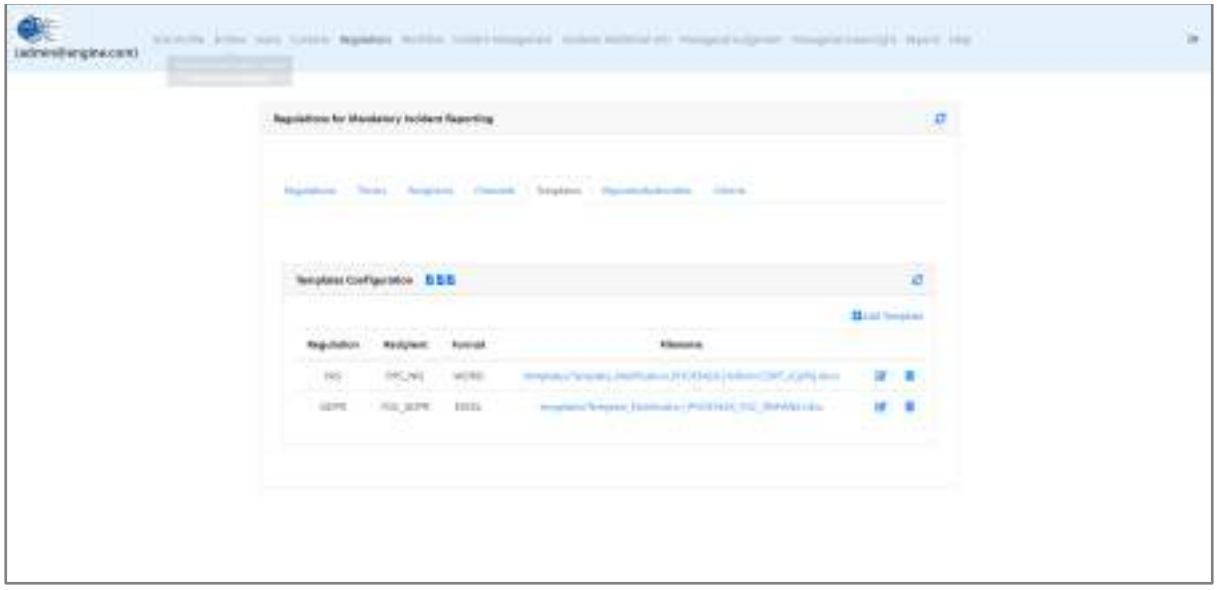

Figure 42: PHOENI2X Report templates configuration in SMIR

5.1.1 Recap of Current Status & Next Steps

- Adapting the tool to the health use case.

5.2 Notification Playbooks

As described in D4.1 [4], PHOENI2X has introduced a new type of playbook in the CACAO standard [1], called the notification playbook. This playbook focuses on orchestrating the steps needed to notify and disseminate information about security events, incidents, or threats. It includes machine-readable metadata for enhanced searching, indexing, and filtering.

During the last reporting period, PHOENI2X has successfully implemented notification playbooks for encoding operating procedures related to incident notifications, CTI forwarding, and incident reporting in the developed use cases.

5.3 Playbook Exchange via MISP & STIX (Information Exchange)

The PHOENI2X consortium is advancing the integration of cybersecurity playbooks with Cyber Threat Intelligence (CTI) through extensions for MISP and STIX 2. These extensions facilitate the sharing of security playbooks, including CACAO, and enable features like semantic indexing and filtering. The design includes a harmonized template for rich metadata that supports various playbook formats. This approach enables organizations to exchange defensive tradecraft in addition to CTI, allowing organizations to respond faster and more effectively to threats.

During the last reporting period, PHOENI2X has updated the schemas with common metadata template that drives the development of the extensions based on the feedback received from end users, advancing the extension to v4. In addition, regarding STIX 2.1, the consortium has performed all the necessary actions required by the OASIS CTI technical committee to prepare and include this extension in the next version of the STIX standard.

The MISP object template can be found on the official GitHub repo[7] of the MISP project. The STIX 2.1 extension and documentation are available on a GitHub repository[8] contributed by University of Oslo to the Open Cybersecurity Alliance, where PHOENI2X project was acknowledged.

Our extensions are agnostic of the playbook format or serialization and can support CACAO or any other open or proprietary format, including playbooks in human natural language or plain graphical representations.

[7] https://github.com/MISP/misp-objects/blob/main/objects/security-playbook/definition.json
[8] https://github.com/opencybersecurityalliance/stix-extensions/blob/main/contexts/playbook/STIX2.1_COA_Playbook_Extension_v4.asciidoc

6 Conclusions

This deliverable provided overview, design & implementation details of the second and final version of Coordinated Response & Preparedness enablers, as delivered in M35 of the project. These included all components developed under the different WP4 tasks, including the Resilience Orchestration & Response editor & engine, the Resilience Playbooks themselves, the Resilience Cyber Range, different types of Serious Games, the Smart Mandatory Incident Reporting tool, the CACAO standard, a specific playbook type for notifications, the layout extension for graphically representing CACAO playbooks, and MISP and STIX 2.1 objects that enhance existing CTI approaches with the ability to exchange defensive tradecraft (i.e., security playbooks).

The output components were leveraged in the context of WP5 activities to update the final version of PHOENI2X, as needed, in order to derive the final release (V2) of PHOENI2X, then carrying out the relevant demonstration and validation activities described in D5.3 - "PHOENi2X framework - Final" and D5.4 - "Evaluation result and PHOENi²X Framework Documentation" by the end of the project in M36.

References

[1] OASIS, Collaborative Automated Course of Action Operations (CACAO) Security Playbooks Version 2.0, Nov. 2023. https://docs.oasis-open.org/cacao/security-playbooks/v2.0/cs01/security-playbooks-v2.0-cs01.html

[2] Somarakis, I., Smyrlis, M., Fysarakis, K., Spanoudakis, G. (2020). Model-Driven Cyber Range Training: A Cyber Security Assurance Perspective. In: Fournaris, A., et al. Computer Security. IOSEC MSTEC FINSEC 2019 2019 2019. Lecture Notes in Computer Science, vol 11981. Springer, Cham. https://doi.org/10.1007/978-3-030-42051-2_12

[3] Smyrlis, M., Fysarakis, K., Spanoudakis, G., Hatzivasilis, G. (2020). Cyber Range Training Programme Specification Through Cyber Threat and Training Preparation Models. In: Hatzivasilis, G., Ioannidis, S. (eds) Model-driven Simulation and Training Environments for Cybersecurity. MSTEC 2020. Lecture Notes in Computer Science(), vol 12512. Springer, Cham. https://doi.org/10.1007/978-3-030-62433-0_2

[4] PHOENI2X D4.1 – Coordinated Response & Preparedness Enablers V1. Konstantinos Fysarakis et al. March 2024

[5] SUNRISE D6.5 – Cyber-physical resilience pilot report V3. Susana González Zarzosa et al. April 2025.