Harmonizing People, Processes, and Technology for Robust Cybersecurity

D3.4 Risk Scenario Analysis Toolkit V2

Document Summary Information

Grant Agreement No	101158555	**Acronym**	CYberSynchrony
Full Title	Harmonizing People, Processes, and Technology for Robust Cybersecurity		
Start Date	01/06/2024	**Duration**	36 months
Project URL	https://www.CYberSynchrony.eu/		
Deliverable	D3.4 Risk Scenario Analysis Toolkit V2		
Work Package	WP3 Comprehensive Threat Intelligence and Customized Risk Evaluation (CYBERRA) and Cybersecurity Resilience and Incident Response Enhancement (CYRESCUE)		
Contractual due date	31/05/2026	**Actual submission date**	07/07/2026
Type	Other	**Dissemination Level**	Public
Lead Beneficiary	MER		
Responsible Author	Andrea Purificato (MER)		
Contributions from	Dimitris Koutras, Malamas Vangelis (FP)		

Funded by the European Union

ECCC

Funded by the European Union under Grant Agreement no. 101158555 and supported by the European Cybersecurity Competence Centre (ECCC). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Cybersecurity Competence Centre. Neither the European Union nor the granting authority can be held responsible for them.

Revision history

Version	Issue Date	Changes	Contributor(s)
0.1	01.05.2026	Initial content	Andrea Purificato (MER), Koutras Dimitris (FP)
0.2	08/05/2026	10% content	Andrea Purificato (MER), Koutras Dimitris (FP)
0.3	25/05/2026	60% content	Andrea Purificato (MER)
0.4	05/06/2026	Update Chapters 2 & 6	Andrea Purificato (MER)
0.5	11/06/2026	Contribution in Chapters 3, 4, 5	Koutras Dimitris (FP)
0.6	12/06/2026	80% content	Andrea Purificato (MER)
0.7	15/06/2026	Updating contribution	Koutras Dimitris (FP)
0.8	18/06/2026	Peer-Review	Konstantinos Adrikos (MAG)
0.9	26/06/2026	Addressing comments from internal review	Andrea Purificato (MER)
0.9b	03/07/2026	QA review	Eleni Keliri (eBOS)
1.0	07/07/2026	Final version for submission	Violeta Vasileva (MAG)

Disclaimer

The content of the document herein is the sole responsibility of the publishers and it does not necessarily represent the views expressed by the European Commission or its services.

While the information contained in the documents is believed to be accurate, the authors(s) or any other participant in the CYberSynchrony consortium make no warranty of any kind with regard to this material including, but not limited to the implied warranties of merchantability and fitness for a particular purpose.

Neither the CYberSynchrony consortium nor any of its members, their officers, employees or agents shall be responsible or liable in negligence or otherwise howsoever in respect of any inaccuracy or omission herein.

Without derogating from the generality of the foregoing neither the CYberSynchrony consortium nor any of its members, their officers, employees or agents shall be liable for any direct or indirect or consequential loss or damage caused by or arising from any information advice or inaccuracy or omission herein.

Copyright message

© CYberSynchrony Consortium This Deliverable contains original unpublished work except where clearly indicated otherwise. Acknowledgement of previously published material and of the work of others has been made through appropriate citation, quotation or both. Reproduction is authorised provided the source is acknowledged.

[1] According to CYberSynchrony's Quality Assurance Process for reports:

- 3 months before Deliverable's Due Date: Table of Contents ready to be reviewed by WP leader/ Technical Manager/ Quality Manager
- 2 months before Deliverable's Due Date: 50% should be complete. Review by the Quality Manager (& Technical Manager for software).
- 1 month before Deliverable's Due Date:100% should be complete. Review by the Quality Manager and Peer Reviewers.
- M: Sent to PC for Submission to the EC after addressing all comments by Quality Manager and Peer Reviewers

Table of Contents

Executive Summary 7
1 Introduction 8
1.1 Mapping CYberSynchrony Outputs 8
1.2 Linkage to other project outputs 9
1.3 Deliverable Structure 10
2 Kitsune Monitoring 11
2.1 Updates since v1 11
2.2 Integration, interoperability and deployment 11
3 Cybersecurity Posture Assessment (CSPA) 13
3.1 Updates since v1 14
3.2 Integration, interoperability and deployment 14
4 Attack Surface Assessment (ASA) 17
4.1 Updates since v1 17
4.2 Integration, interoperability and deployment 18
5 Purple Team Assessment (PTA) 21
5.1 Updates since v1 22
5.2 Integration, interoperability and deployment 22
6 KPIs Mapping 25
6.1 KPI3.7 Risk Mitigation Index 25
7 Conclusions 28
8 References 29
Annex I: PoC and Prototypes (Kitsune Monitoring) 30
Annex II: Integration of ASA, CSPA into the CYberSynchrony Platform 33

List of Figures

Figure 1: CSPA Architecture 13
Figure 2: CSPA End-to-End Flow 16
Figure 3: ASA Architecture 17
Figure 4: ASA End-to-End Flow 20
Figure 5: PTA Architecture 21
Figure 6: PTA End-to-End Flow 24
Figure 7: Kitsune Login 30
Figure 8: Kitsune Homepage 31
Figure 9: Kitsune Incidents Overview 32
Figure 10: Kitsune Incidents Details 32
Figure 11: ASA login access 33
Figure 12 : ASA Detailed View: Operational Posture Mapping 34
Figure 13: CSPA login access 35
Figure 14: CSPA Detailed View: Operational Posture Mapping 35

List of Tables

Table 1: Adherence to CYberSynchrony's GA 8
Table 2: Links to other Project Outputs/Work 9
Table 3: CSPA Layer / Responsibility 15
Table 4: CSPA Layer / Meaning 15
Table 5: ASA Layer / Responsibility 18
Table 6: PTA Layer / Responsibility 23

Acronyms and abbreviations

Abbreviation / Term	Description
API	Application Programming Interface
APT	Advanced Persistent Threat
ASA	Attack Surface Assessment
ATT	Adversarial Tactics, Techniques
CAPEC	Common Attack Pattern Enumeration and Classification
CPE	Common Platform Enumeration
CPU	Central Processing Unit
CRUD	Create, Read, Update, Delete
CSPA	Cybersecurity Posture Assessment
CSV	Comma-Separated Values
CTI	Cyber Threat Intelligence
CVE	Common Vulnerabilities and Exposures
CVSS	Common Vulnerability Scoring System
CWE	Common Weakness Enumeration
DAO	Data Access Object
DNS	Domain Name System
DPM	Data Preparation Module
GDPR	General Data Protection Regulation
GUI	Graphical User Interface
HEUR	Heuristic
HTTP	Hypertext Transfer Protocol
IDS	Intrusion Detection System
IPS	Intrusion Prevention System
JSON	JavaScript Object Notation
KPI	Key Performance Indicator
LAMP	Linux, Apache, MySQL, PHP (web stack)

LDAP	Lightweight Directory Access Protocol
LEMP	Linux, Nginx (Engine-X), MySQL, PHP (web stack)
LLC	Limited Liability Company
LLM	Large Language Model
MFA	Multi-Factor Authentication
MISP	Malware Information Sharing Platform
MITRE	MITRE Corporation
MS	Member States
NIST	National Institute of Standards and Technology
NVD	National Vulnerability Database
PHP	Hypertext Preprocessor
PoC	Proof of Concept
PTA	Purple Team Assessment
RAW	Raw data format
RBAC	Role-Based Access Control
REST	Representational State Transfer
SAML	Security Assertion Markup Language
SDK	Software Development Kit
SHA	Secure Hash Algorithm
SOC	Security Operations Centre
STIX	Structured Threat Information Expression
SVG	Scalable Vector Graphics
TLP	Traffic Light Protocol
UFT	Unalterable Forensic Techniques
WHOIS	Domain Registry Lookup Protocol
XML	eXtensible Markup Language
YAML	YAML Ain't Markup Language

Executive Summary

Deliverable D3.4, titled "*Risk Scenario Analysis Toolkit V2*", falls under WP3 of the CYberSynchrony Project; and it is the consolidated and production-ready version of the software suite, which updates the initial framework established in D3.2 by integrating dynamic risk parameter refinement and real-time threat data to ensure more adaptive and technically mature risk evaluations.

This toolkit integrates four key standalone components, *Kitsune Monitoring*, *Cybersecurity Posture Assess*ment (CSPA), *Attack Surface Assessment* (ASA), and *Purple Team Assessment* (PTA) to provide a holistic, evidence-based approach to cyber threat evaluation and incident preparedness.

- Kitsune Monitoring offers proactive surveillance for early detection of threats across open and clandestine sources.
- The Cybersecurity Posture Assessment (CSPA) delivers a control-mapped, real-world assessment of organizational security maturity, grounded in operational data and aligned with industry standards.
- The Attack Surface Assessment (ASA)automates endpoint discovery and vulnerability analysis, enabling rapid identification and remediation of risks.
- The Purple Team Assessment (PTA) bridges offensive and defensive operations, simulating adversary behaviours to validate and enhance detection and response capabilities.

Building upon the foundational capabilities presented, D3.4, titled "*Risk Scenario Analysis Toolkit v2*" will provide an enhanced and extended version of the D3.2 "*Risk Scenario Analysis Toolkit v1*"; this updated release incorporates refinements based on stakeholder feedback, operational testing, and evolving threat landscapes.

1 Introduction

This section presents the CYberSynchrony's Grant Agreement (GA) commitments, as extracted from the formal deliverable and task description, with respect to their outputs and work to be performed.

1.1 Mapping CYberSynchrony Outputs

Table 1 below provides a mapping of D3.4 sections to related deliverables and describes the specific reference or dependency for each, ensuring project-wide alignment and awareness of cross-WP interactions.

Table 1: Adherence to CYberSynchrony's GA

Component no	Component title	Respective Document Chapter(s)	Justification
DELIVERABLE			
D3.4	Risk Scenario Analysis Toolkit V2	Chapter 1,7	Defining a software suite offering tools for evaluating and understanding specific cyber risk scenarios.
		Chapter 2 (Kitsune Monitoring), Annex I	In the context of WP3, Kitsune Monitoring component contributes to the strengthening of risk monitoring services and the early identification of cyber incidents, offering specialized risk scenario analysis tailored for MS.
		Chapter 3 (CSPA), Annex II	The CSPA service contributes directly to project's objective of supporting customized, operational cyber risk assessments for MS.
		Chapter 4 (ASA), Annex II	The ASA component is closely aligned with the overarching objectives of the project, as it provides a proactive, context-aware evaluation of cyber risks tailored to each MS. By simulating real adversary behaviours, customizing attack vectors to reflect national infrastructure and threat profiles, and identifying exploitable attack chains, ASA offers a realistic and operationally relevant stress test of existing defences.
		Chapter 5 (PTA), Annex II	PTA service by simulating realistic attack scenarios and assessing how well current security measures detect and respond to them, helps identify vulnerabilities and validate the effectiveness of detection mechanisms. This dual focus on exploiting weaknesses and strengthening response capabilities aligns with WP3's aim to

			enhance real-time threat evaluation and improve incident preparedness, ultimately empowering organizations and MS with more resilient and adaptive cybersecurity postures.
TASKS			
T3.3	Crafting Personalized Risk Assessment Tools	Chapter 2,3,4,5, Annex I, II	This task aims to develop a versatile software framework tailored to the diverse cybersecurity landscapes of MS. The framework excels in creating customized risk assessment metrics based on each state's specific threats and vulnerabilities.
T3.4	Dynamic Risk Parameter Refinement	Chapter 3,4,5, Annex II	By continuously feeding real-time threat data into its risk parameter engine, the framework ensures that its threat models are always in sync with the most current cyber threat landscapes. This relentless adaptation equips MS with risk profiles that remain perennially relevant, thereby enabling them to remain agile in the face of an ever-transforming cyber environment.

1.2 Linkage to other project outputs

The table below explains how this deliverable is related to other project outputs and specific deliverables.

Table 2: Links to other Project Outputs/Work

WP Number	Task Number	Deliverable Number related	Content
WP2	T,2.1, T2.3, T2.4	D2.1	Deliverable D2.1 serves as a foundational element for the overall architecture of the CYberSynchrony ecosystem, making it highly relevant to D3.4.
WP3	T.3.4	D3.1, D3.3	For D3.4 to effectively perform its risk analysis, its specifications, such as data input formats, expected types of threat indicators, and analytical capabilities, must be aligned with and informed by the outputs and capabilities of D3.1 and D3.3.
		D3.2	D3.4 updates the initial framework established in D3.2
WP4	T4.3, T4.4	D4.1	D4.1 provides the crucial monitoring capabilities necessary to assess D3.4's performance and

			functional integration within the broader CYberSynchrony ecosystem.
WP8	T8.1	D8.1	D8.1 serves as the foundational document detailing the integration roadmap for all CYberSynchrony modules, including components from Work Package 3 (WP3) such as D3.4.

1.3 Deliverable Structure

This document is divided into the following chapters:

- **Chapter 1: Introduction**

This chapter is an overview of Deliverable 3.4. It introduces us to the Document's objectives and its contribution to the overall project while it outlines its alignment with it.

- **Chapter 2: Kitsune Monitoring**

This chapter focuses on one of the four tools defining the Toolkit, Kitsune. It highlights the component's purpose and how it aligns with the objective of the Project and WP. In addition to this, it provides highlights on tool's architecture, functionalities, inner workings and its advantages over existing solutions while showcasing its programming languages, frameworks and libraries.

- **Chapter 3: Cybersecurity Posture Assessment (CSPA)**

This chapter goes deep on analysing Cybersecurity Posture Assessment, the second tool to be used for Risk Scenario Analysis in the Ecosystem of CYberSynchrony. It provides useful information regarding the tool's purpose and its alignment with the CYberSynchrony Project while bringing knowledge on component's architecture and functionalities as well as on its technology.

- **Chapter 4: Attack Surface Assessment (ASA)**

This chapter presents Attack Surface Assessment as one of the four components of the Risk Scenario Analysis Toolkit. It focuses on tool's purpose and its alignment with the objective of the Project and WP. It goes further than this, describing tool's architecture, functionalities, inner workings and its advantages over existing solutions, offering at the same time highlights on its programming languages, frameworks and libraries.

- **Chapter 5: Purple Team Assessment (PTA)**

This chapter discusses Purple Team Assessment as the last of the four tools presented in this Deliverable. As in the previous chapters, the main focus is on the purpose of the component and its alignment with the objective of the Project and the WP. It includes highlights on tool's architecture, functionalities, inner workings and its advantages over existing solutions. At the same time, it underlines the programming languages, frameworks and libraries used.

- **Chapter 6: KPIs Mapping**

This chapter maps monitoring suite components documented in this deliverable to the project KPIs to which they contribute.

- **Chapter 7: Conclusion**

This concluding chapter summarizes the output of Risk Scenario Analysis Toolkit while analyzing its four components, their role in defining the Risk Scenario Analysis Toolkit V2 and their contribution to the whole CYberSynchrony Ecosystem.

2 Kitsune Monitoring

While the foundational version of Kitsune Monitor focused on providing backend RESTful APIs for machine-to-machine communication, version 2 introduces a web-based front-end integrated into the CYberSynchrony platform.

This interface allows SOC analysts and operators to move beyond raw JSON data and interact with threat intelligence through intuitive dashboards and visual analytics.

Once logged in, the user can navigate through several specialized views that transform Kitsune's extensive database of threat actors, botnets, and data leaks into actionable security insights.

This technical integration is fundamental to addressing and achieving the project's Key Performance Indicators:

- **KPI 1.1** (Number of Threat Assessments Carried Out): by automating the flow of dark web alerts into the platform's unified view, Kitsune Monitoring ensures that every identified risk related to monitored keywords is captured in the Threat Assessment Register, contributing to the project's requirement of >36 documented assessments.
- **KPI 3.7** (Risk Mitigation Index): the real-time context provided by Kitsune Monitoring enables analysts to perform more accurate "what if" scenario evaluations. This data-driven approach directly supports the target of a 20% reduction in identified risks by identifying exposures before they can be exploited in a real-world incident.
- **KPI 5.1** (Number of Risk Monitoring Services Provided): the consolidated deployment of Kitsune Monitoring profiles across the healthcare, public administration, and academic pilots is counted towards the project benchmark of activating 25 risk monitoring services. This integration ensures that these services provide the required accuracy and timely alerts necessary to maintain the project's target of a >95% user satisfaction rating.

2.1 Updates since v1

For deliverable D3.4 (Risk Scenario Analysis Toolkit v2), which updates the initial framework established in D3.2, the following paragraph describes the integration of Kitsune Monitoring and its access via the centralized platform's interface.

While the technical characteristics, endpoints, and specifications of the RESTful APIs remain comprehensively detailed in the first version of this deliverable (D3.2), this updated version (D3.4) integrates the specific work performed to achieve integration and interoperability within the CYberSynchrony platform.

A major achievement in this phase is the deployment of a web-based graphical user interface (GUI) integrated into the project's centralized dashboard. This front-end provides authorized analysts with direct access to Kitsune services, transforming backend JSON intelligence into intuitive visual modules for tracking threat actors, managing the digital attack surface, and triaging underground information leaks.

By embedding these capabilities into the broader modular architecture, the toolkit ensures that proactive monitoring data, such as botnet infections or dark web mentions, is presented in a human-readable and actionable format, significantly enhancing real-time threat detection and adaptive risk evaluation workflows.

Additional details can be found in the Annex I: PoC and Prototypes (Kitsune Monitoring) section.

2.2 Integration, interoperability and deployment

In v2 of the *Risk Scenario Analysis Toolkit (D3.4)*, the Kitsune Monitoring component evolves from the foundational framework established in v1 into a fully production-ready service available in the CYberSynchrony modular architecture.

Interoperability is achieved through standard RESTful APIs exchanging data in JSON format, enabling the seamless delivery of security alerts directly to the centralized CYBERRA module dashboard and a web front-end within the unified CYberSynchrony platform architecture.

This integration allows the platform to automatically correlate Kitsune's proactive intelligence, such as botnet infections, dark web mentions, and data leak exposures, with internal telemetry from Wazuh[1] and global feeds from MISP, providing a multi-layered enrichment of the threat landscape for pilot organizations.

3 Cybersecurity Posture Assessment (CSPA)

The Cybersecurity Posture Assessment (CSPA) component is the posture-mapping capability of the CYRESCUE module within the CYberSynchrony platform. In the context of D3.4, CSPA represents the evolved and operational version of the posture assessment functionality initially introduced in D3.2. While the first version described the conceptual direction and preliminary prototype of the component, the current version provides a more concrete, integrated, and technically mature tool that supports asset-based security posture analysis through automated inventory ingestion, asset normalization, CPE resolution, and mapping against the MITRE D3FEND[2] and MITRE ATT&CK[3] knowledge domains.

The main purpose of CSPA is to support organisations in understanding the defensive posture associated with their digital assets. It does so by receiving a structured host inventory, identifying the hardware, software, and services present in the environment, and correlating them with defensive knowledge models. Through this process, CSPA allows stakeholders to understand which types of digital artifacts are present, which adversarial techniques may target them, which detection strategies are applicable, and which mitigations can be considered. This supports the broader CYberSynchrony objective of enabling evidence-based risk evaluation and improving preparedness against cyber threats.

CSPA operates on inventory data provided in JSON format. The inventory includes three main artifact categories: hardware, applications, and services. It may also include host-level information such as hostname, operating system, operating system version, domain, uptime, manufacturer, and model. The dashboard forwards the uploaded inventory to the backend without pre-processing it, using an application/octet-stream request. The backend is then responsible for parsing the inventory, normalising vendor and product strings, resolving assets to Common Platform Enumeration (CPE) identifiers, and linking the resulting entities to the posture knowledge graph.

A key aspect of CSPA is the correlation of each identified asset with the MITRE D3FEND ontology. For each matched artifact, the system identifies the relevant DigitalArtifactType, the MITRE ATT&CK techniques that may target that artifact, the detection strategies that can be used to identify related adversarial behaviour, and the mitigations that can reduce or prevent the corresponding risk. This enables the component to move beyond a simple asset inventory and provide an actionable view of the defensive landscape surrounding each asset.

From a user perspective, CSPA provides an assessment summary dashboard that offers high-level situational awareness. The dashboard displays the total number of discovered artifacts, a per-type breakdown of hardware, applications, and services, and a host context card presenting the main host metadata. The interface also provides colour-coded chips and tabbed navigation, allowing the user to review each resolved artifact separately. Within each artifact tab, the user can examine the matched product, vendor, version, CPE string, Digital Artifact Ontology classification, and confidence score of the CPE match. The D3FEND and ATT&CK views are rendered through nested expandable panels, allowing the defensive information to remain readable even when a single asset is associated with multiple techniques, detections, or mitigations.

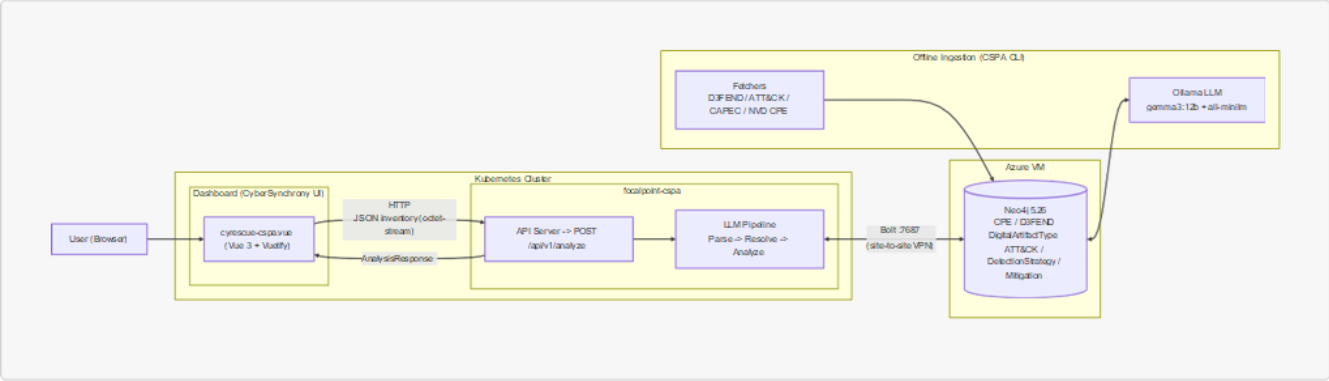

Figure 1: CSPA Architecture

3.1 Updates since v1

Compared with the first version of the Risk Scenario Analysis Toolkit, the CSPA component has evolved from a preliminary posture-assessment concept into a more complete and integrated service. The most important update is the transition from a general description of cybersecurity posture assessment to a concrete asset-to-posture correlation workflow. The current version supports the ingestion of structured host inventories, automatic normalisation of vendor and product names, resolution of assets to CPE identifiers, and mapping of the resulting artifacts to a pre-built posture graph.

The updated CSPA implementation introduces a stronger knowledge-driven approach. Instead of only presenting assessment results at a high level, the component now links each asset to the relevant MITRE D3FEND DigitalArtifactType nodes, associated MITRE ATT&CK techniques, detection strategies, and mitigations. This enables a more detailed and operationally useful assessment, as users can directly observe how specific assets relate to possible adversarial behaviours and defensive measures.

Another important update concerns the user interface. The CSPA dashboard has been refined to support inventory upload, summary presentation, artifact-level navigation, and nested defensive analysis. The tabbed artifact view provides a clear structure for reviewing each matched asset, while the expandable D3FEND and ATT&CK panels allow users to explore deeper contextual information without overloading the interface. This makes the tool more usable for both technical analysts and stakeholders who require a structured overview of the security posture.

The backend has also been enhanced. The CSPA service is implemented as a Go-based backend using the Chi router and Huma v2 for OpenAPI-driven endpoint registration. The processing pipeline follows a typed multi-stage flow, including parsing, CPE resolution, and analysis enrichment. Parallel worker-pool batching is used to process inventories efficiently. The component uses Neo4j as its graph database, where the posture graph is stored and queried at runtime. The graph includes CPEs, DigitalArtifactType entities, ATT&CK techniques, detection strategies, and mitigations.

The component further integrates an LLM-assisted enrichment layer. Ollama is used to serve the gemma3:12b model for CPE DAO classification and technical-detail extraction, while all-minilm is used for embeddings and retrieval-augmented processes. These capabilities support the enrichment of technical metadata and improve the quality of asset classification during data preparation.

The current version also introduces a clearer deployment model. CSPA is containerised using a multi-stage Dockerfile and deployed in a Kubernetes environment. The service runs as a non-root container, uses a read-only root filesystem, drops unnecessary capabilities, and exposes health checks for liveness and readiness. Network access is restricted through a NetworkPolicy, and configuration is managed through ConfigMaps and Secrets. This deployment approach is aligned with secure-by-design principles and supports integration into the wider CYberSynchrony platform environment.

3.2 Integration, interoperability and deployment

CSPA is integrated into the CYberSynchrony platform as part of the CYRESCUE module. The frontend is implemented as a Vue 3 and TypeScript page using Vuetify 3 components. The relevant dashboard page handles file selection, upload, request dispatch, and rendering of the returned analysis results. It provides the user-facing layer through which operators can upload inventory files and explore posture analysis outcomes.

The backend is provided by the focalpoint-cspa service. It receives the uploaded inventory through a REST endpoint and assigns a correlation identifier for traceability. The service validates the request, parses the inventory, normalises vendor and product values, resolves artifacts to CPE nodes, and traverses the posture graph to retrieve the relevant D3FEND and ATT&CK information. The returned response includes status information, correlation metadata, processing duration, matched artifacts, confidence scores, and the corresponding defensive mappings.

The CSPA architecture consists of several cooperating layers. The dashboard layer manages user interaction and result rendering. The API layer provides ingress, request validation, OpenAPI documentation, and correlation tracking. The LLM pipeline performs typed and parallel processing of the inventory payload. The Neo4j graph database stores the pre-built posture graph, linking CPEs with DigitalArtifactType nodes, attack techniques, detection strategies, and mitigations. The Ollama engine enriches CPE nodes with DAO classifications and technical metadata during data preparation. Finally, periodic fetchers refresh the underlying knowledge sources, including MITRE D3FEND, MITRE ATT&CK STIX, CAPEC, and NVD[4] CPE data.

Table 3: CSPA Layer / Responsibility

Layer	Responsibility
Dashboard (Vue)	File selection, upload, nested accordion rendering of D3FEND / ATT&CK / Detection / Mitigation layers.
CSPA API	Auth-less ingress, correlation-ID tracking, OpenAPI docs, request validation.
LLM Pipeline	Typed, parallel processing of an inventory payload — Parse → CPE Resolve → Analysis Enrichment.
Neo4j graph DB	Holds the pre-built posture graph: CPEs ↔ DigitalArtifactType ↔ AttackTechnique ↔
Ollama Engine	DetectionStrategy / Mitigation.
Fetchers	Enriches CPE nodes with DAO classification and technical metadata during data preparation.

The component is deployed as a containerised service within Kubernetes. The container image is published under the CYberSynchrony GitLab registry [5] .The Kubernetes deployment uses a non-root runtime user, read-only root filesystem, dropped Linux capabilities, and health probes on the /health endpoint. The service is exposed internally through a ClusterIP on port 8445, while ingress is restricted to TCP/8445 through a NetworkPolicy. Runtime configuration, including Neo4j connection details and service parameters, is provided through ConfigMaps and Secrets. The standard resource footprint is defined as 100m CPU and 128Mi memory requested, with limits of 500m CPU and 512Mi memory per pod.

From an interoperability perspective, CSPA relies on widely adopted cybersecurity data models and standards, including CPE for platform identification, MITRE D3FEND for defensive ontology mapping, MITRE ATT&CK for adversarial technique mapping, and CAPEC for attack-pattern context. The use of structured JSON input, RESTful APIs, OpenAPI-driven endpoint documentation, and graph-based knowledge representation ensures that the component can interoperate with other CYberSynchrony modules and external cybersecurity tooling.

Table 4: CSPA Layer / Meaning

Layer	Meaning
Digital Artifact	D3FEND DigitalArtifactType node (id, label, definition) — what the asset is, ontologically.

ATT&CK Technique	Adversary technique that targets that Digital Artifact (technique ID such as T1021, T1190).
Detection Strategy	D3FEND analytic that detects the technique (id + human-readable name).
Mitigation	DetectionStrategy / Mitigation.
Fetchers	D3FEND mitigation that prevents/reduces the technique (id, name, full description).

The end-to-end user-facing sequence starts when the user uploads an inventory JSON file through the CYberSynchrony dashboard and selects the analysis action. The dashboard sends the file to the focalpoint-cspa backend through a POST request. The backend parses the host and artifact information, normalises vendor and product names, resolves assets to CPE nodes, and enriches the results by traversing the Neo4j posture graph. The response is then returned to the dashboard, where the summary bar, artifact tabs, and expandable D3FEND / ATT&CK panels are rendered for analyst review.

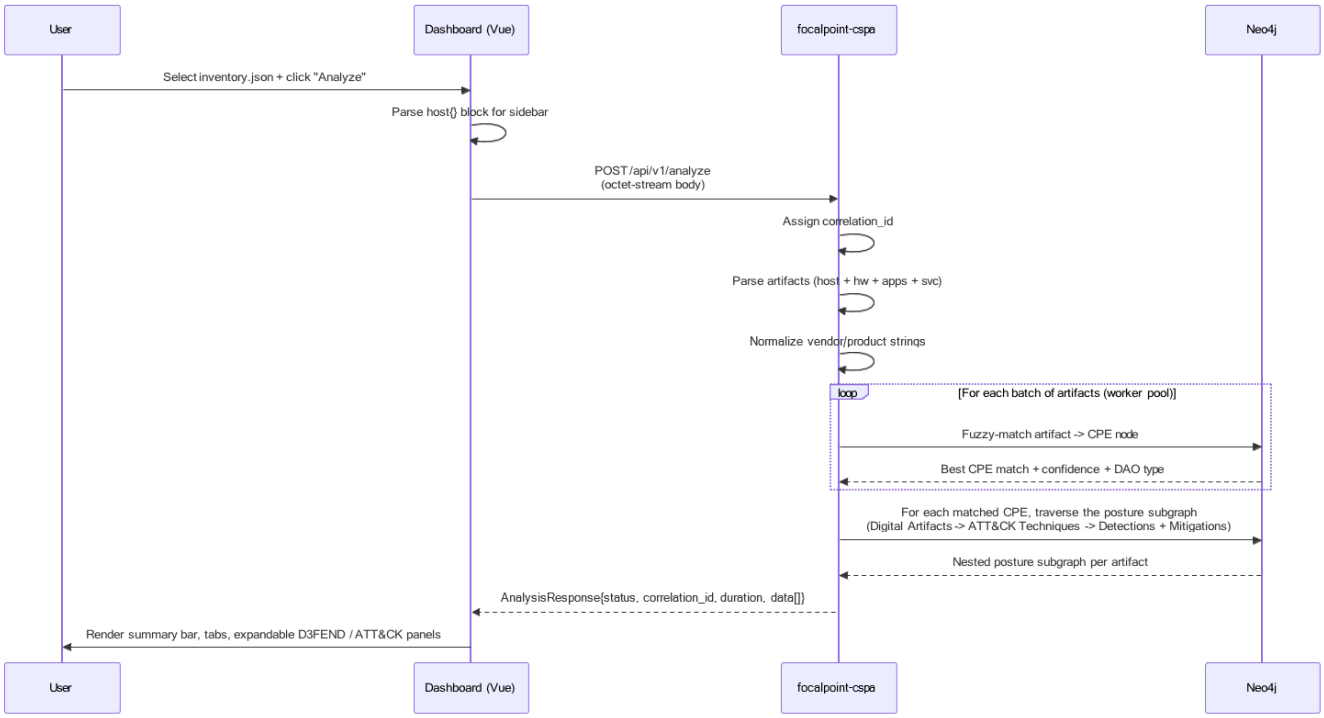

Figure 2: CSPA End-to-End Flow

4 Attack Surface Assessment (ASA)

The Attack Surface Assessment (ASA) component is the vulnerability-discovery capability of the CYBERGOPLUS module within the CYberSynchrony platform. In D3.4, ASA represents the evolved version of the attack surface assessment functionality initially described in D3.2. Its role is to support the identification, correlation, and prioritisation of vulnerabilities affecting assets discovered in an organisation's environment.

ASA receives a structured host inventory containing hardware, applications, and services, optionally enriched with host metadata such as hostname, operating system, operating system version, domain, uptime, manufacturer, and model. The component resolves each discovered asset to a normalised CPE identifier and correlates the identified CPEs against curated vulnerability data, including CVEs[6] and CWEs[7], sourced from the NIST National Vulnerability Database and MITRE knowledge bases. Through this process, ASA enables users to identify vulnerable technologies, understand associated weaknesses, and prioritise remediation based on severity and exposure.

The component is designed to provide a practical view of the attack surface. Rather than operating as a generic scanner, ASA focuses on correlating asset inventory information with authoritative vulnerability and weakness datasets. The backend normalises vendor and product strings before matching them to CPEs, allowing inconsistent or verbose product names to be mapped to standardised identifiers. Once an asset is resolved, the system queries the graph database to identify associated vulnerabilities and weakness patterns.

The ASA dashboard provides a high-level assessment summary, including the total number of parsed artifacts, per-type breakdown of hardware, applications, and services, and host context information. Colour-coded chips are used to present artifact counts and contextual indicators. The component also includes a tabbed artifact view, where each resolved artifact can be reviewed individually. For every artifact, the user can inspect the matched product, vendor, version, CPE string, DAO classification, and confidence score for the CPE match.

A central feature of ASA is the artifact-level CVE view. For each matched artifact, the system displays the associated CVEs together with relevant vulnerability information. This includes the CVE identifier, severity, CVSS base score, publication date, and associated CWE tags. Severity is displayed using colour-coded badges, allowing analysts to quickly identify critical and high-risk vulnerabilities. The interface uses a scrollable grid to remain usable even when a single artifact is associated with a large number of vulnerabilities.

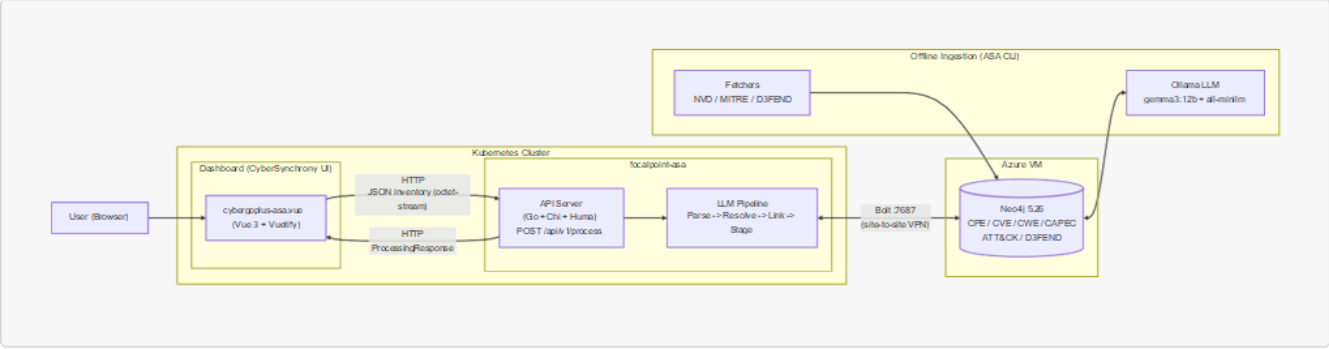

Figure 3: ASA Architecture

4.1 Updates since v1

Since the first version of the Risk Scenario Analysis Toolkit, ASA has evolved from a preliminary attack surface assessment concept into a more technically complete vulnerability-correlation service. The updated version supports structured inventory ingestion, asset normalisation, CPE resolution, and correlation with CVE and CWE data. This provides a more operationally useful output, allowing users to move from raw asset information to vulnerability intelligence in a structured workflow.

One of the main updates is the introduction of a clear asset-to-vulnerability correlation process. The backend parses the uploaded inventory, extracts host and artifact information, normalises vendor and product names, resolves artifacts to CPE identifiers, and links those identifiers to vulnerability information stored in the graph database. This enables the component to identify vulnerabilities associated with discovered assets and present them in a structured and analyst-friendly form.

The current version also introduces a clearer dashboard experience. Users can upload an inventory file, review an assessment summary, navigate between resolved artifacts, and examine associated CVEs. The CVE view includes severity, CVSS score, publication date, and CWE mapping, enabling faster triage and prioritisation. The dashboard also supports both live and demo modes. The live mode sends the uploaded inventory to the focalpoint-asa backend, while the demo mode uses static sample data and does not require backend availability. This supports both operational deployment and demonstration scenarios.

The backend has been implemented as a Go service using the Chi router and Huma v2 for OpenAPI-driven endpoint registration. The LLM pipeline follows a typed multi-stage processing flow, including parsing, resolving, linking, and staging. Parallel worker-pool batching is used to support efficient processing of multiple artifacts. Neo4j is used as the graph database, storing the pre-built knowledge graph of CPEs, CVEs, CWEs, vendor and product nodes, and their relationships.

The component also uses LLM-assisted enrichment. Ollama serves the gemma3:12b model for CPE DAO classification and technical-detail extraction, while all-minilm supports embeddings and retrieval-augmented processes. This strengthens the quality of asset classification and supports richer technical metadata during data preparation.

From a deployment perspective, ASA has been packaged as a containerised service and prepared for Kubernetes deployment. It runs as a non-root container, uses a read-only root filesystem, drops unnecessary capabilities, exposes health probes, and is restricted through a NetworkPolicy. This reflects a significant maturation compared with the preliminary prototype stage and makes the component suitable for integration into the operational CYberSynchrony environment.

4.2 Integration, interoperability and deployment

ASA is integrated into the CYberSynchrony platform as part of the CYBERGOPLUS module. The frontend is implemented using Vue 3, TypeScript, Vite, and Vuetify 3. The relevant dashboard page handles file upload, request dispatch, response rendering, severity styling, and artifact navigation. The user interacts with the component through the CYberSynchrony dashboard, where inventory files can be submitted for processing and the returned vulnerability information can be explored.

The backend is provided by the focalpoint-asa service. It receives the uploaded inventory through a REST endpoint, assigns a correlation identifier, validates the request, parses the inventory, normalises the extracted artifact data, resolves artifacts to CPE nodes, and queries the graph database for associated vulnerabilities. The response includes processing status, correlation metadata, processing duration, matched artifacts, CPE confidence values, and linked CVE information.

The ASA architecture is composed of multiple layers. The dashboard layer manages file selection, upload, result rendering, severity styling, and navigation. The API layer provides request handling, OpenAPI documentation, validation, and correlation tracking. The LLM pipeline performs typed and parallel processing of the inventory payload. The Neo4j graph database stores the pre-built vulnerability knowledge graph, including CPEs, CVEs, CWEs, vendor and product nodes, and relationships. The Ollama engine enriches CPE nodes with DAO classification and technical metadata during data preparation. Fetchers periodically refresh the vulnerability and cybersecurity knowledge sources, including NVD CVE, NVD CPE, MITRE CWE, MITRE CAPEC, MITRE ATT&CK, and MITRE D3FEND data.

Table 5: ASA Layer / Responsibility

Layer	Responsibility
Dashboard (Vue)	File selection, upload, result rendering, severity styling, navigation.
ASA API	Auth-less ingress, correlation-ID tracking, OpenAPI docs, request validation.
LLM Pipeline	Typed, parallel, multi-stage processing of an inventory payload.
Neo4j graph DB	Holds the pre-built knowledge graph: CPEs, CVEs, CWEs, vendor/product nodes and relationships.
Ollama Engine	Enriches CPE nodes with DAO classification and technical metadata during data preparation.
Fetchers	Periodically refresh CVE/CPE/CWE/CAPEC/D3FEND/ATT&CK feeds into Neo4j.

The component is deployed as a containerised service in Kubernetes. The container image is published under the CYberSynchrony GitLab registry. The Kubernetes deployment uses a non-root user, read-only root filesystem, dropped capabilities, and liveness/readiness probes on the /health endpoint. The service is exposed internally as a ClusterIP on port 8444, while ingress is restricted to TCP/8444 through a NetworkPolicy. Neo4j credentials and runtime configuration are provided through ConfigMaps and Secrets. The defined resource footprint is 100m CPU and 128Mi memory requested, with limits of 500m CPU and 512Mi memory per pod.

ASA supports interoperability through the use of standard cybersecurity data models and sources. CPE is used for asset identification, CVE for vulnerabilities, CWE for weakness classification, CAPEC for attack-pattern context, and MITRE ATT&CK and D3FEND for additional adversarial and defensive knowledge. This enables ASA to support structured vulnerability discovery and to interoperate with the wider CYberSynchrony risk analysis ecosystem.

The end-to-end user-facing flow begins when the user uploads an inventory JSON file and selects the processing action in the dashboard. The dashboard sends the inventory to the focalpoint-asa backend. The backend parses the host block and artifact lists, normalises vendor and product names, fuzzy-matches artifacts to CPE nodes, links artifacts to the host node where applicable, queries CVEs for the matched CPEs, and returns the processing response. The dashboard then renders the summary bar, artifact tabs, and CVE tables for analyst review.

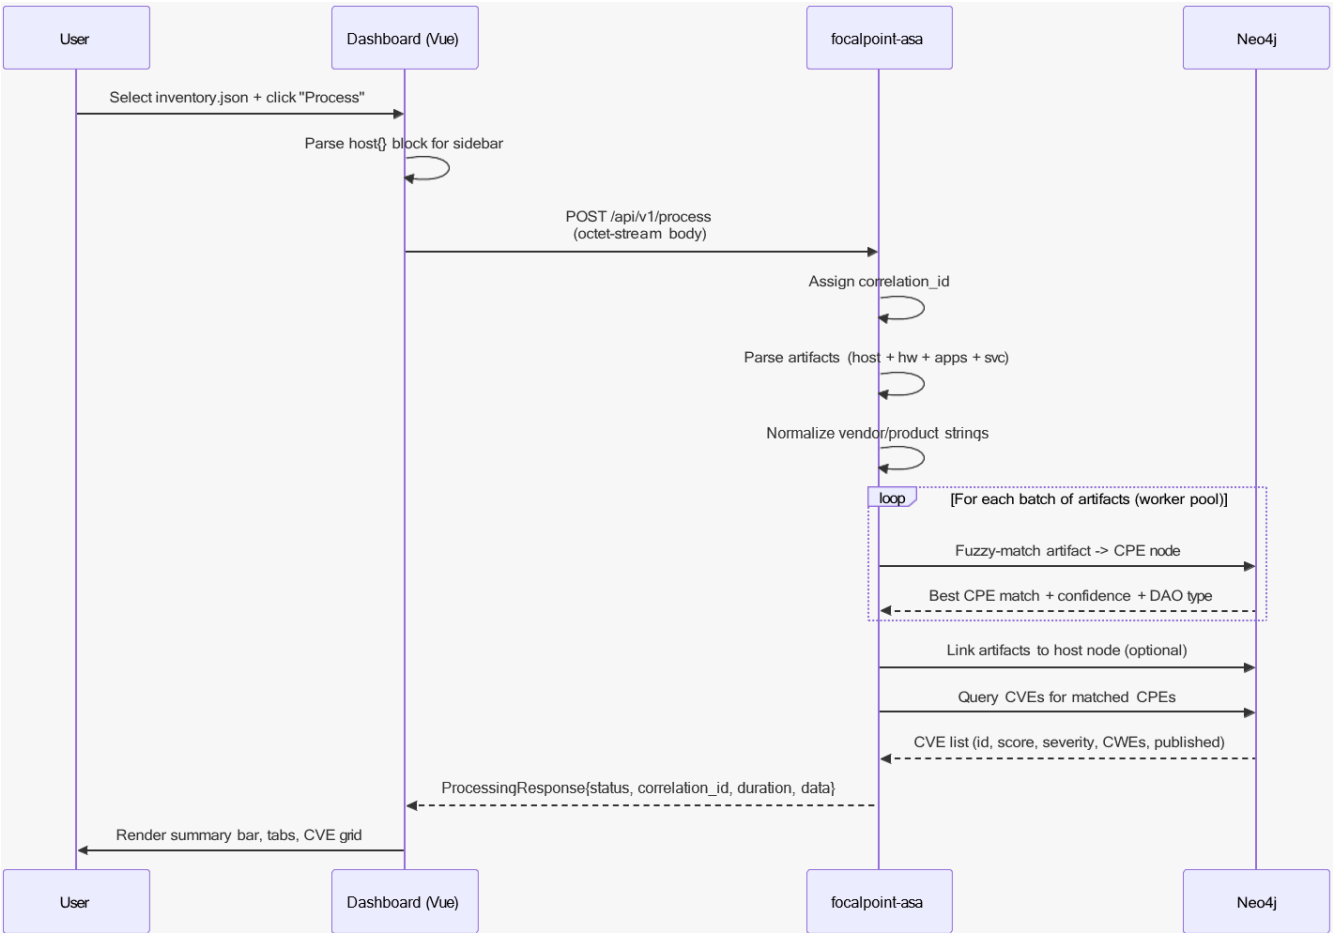

Figure 4: ASA End-to-End Flow

5 Purple Team Assessment (PTA)

The Purple Team Assessment (PTA) component is the combined red-team and blue-team exercise-management capability of the CYRESCUE module within the CYberSynchrony platform. In D3.4, PTA represents the operational version of the purple teaming functionality that supports structured assessment of detection, prevention, logging, and response capabilities against adversarial techniques mapped to MITRE ATT&CK.

The purpose of PTA is to enable operators to transform offensive testing material into structured test cases and to record the defensive performance of an organisation against those test cases. The component supports the import of red-team outputs, such as Caldera[8] reports, MITRE ATT&CK Navigator layers, selected Atomic Red Team templates, and prior PurpleOps campaigns. These inputs are transformed into structured test cases that can be scored by the blue team using outcomes such as Prevented, Alerted, Logged, or Missed. Evidence can be attached to each test case, and results can be exported in several formats for reporting and follow-up.

PTA provides a complete assessment-management workflow. Users can create new assessments, import test cases, organise sources, targets, tools, controls, and tags, edit test-case metadata, record red-team and blue-team notes, upload evidence, review knowledge-base information, inspect related Sigma detection rules, and export the results. This supports the CYberSynchrony objective of improving preparedness and response through practical validation of detection and mitigation capabilities.

The component includes an embedded content corpus compiled into the backend binary, allowing it to operate without runtime network access for its core knowledge sources. This content includes the MITRE ATT&CK catalogue, Atomic Red Team test-case templates, Sigma rules linked to techniques, and knowledge-base entries per technique. A content snapshot is displayed in the dashboard to inform users about the bundled knowledge corpus available in the backend.

The dashboard includes an assessments panel, a templates browser, an assessment detail dialog, and a nested test-case editor. The assessments panel lists existing assessments and allows users to create, open, import, or delete assessments. The templates browser allows users to search the seeded Atomic Red Team library using server-side and client-side filters and import selected templates into an assessment. The assessment detail dialog provides five main tabs: Test cases, Stats, Coverage, Multi-entities, and Imports. These tabs provide a complete working surface for executing, scoring, analysing, and documenting purple-team exercises.

The Test cases tab presents the individual cases, including name, MITRE ID, tactic, state, outcome, and modification status. The Stats tab presents outcome chips and per-tactic breakdowns. The Coverage tab renders a server-generated SVG hexagon heatmap covering the MITRE tactics and scoring them according to defensive outcomes. The Multi-entities tab allows the editing of sources, targets, tools, controls, and tags. The Imports tab provides upload surfaces for Caldera reports, MITRE Navigator layers, and PurpleOps campaigns.

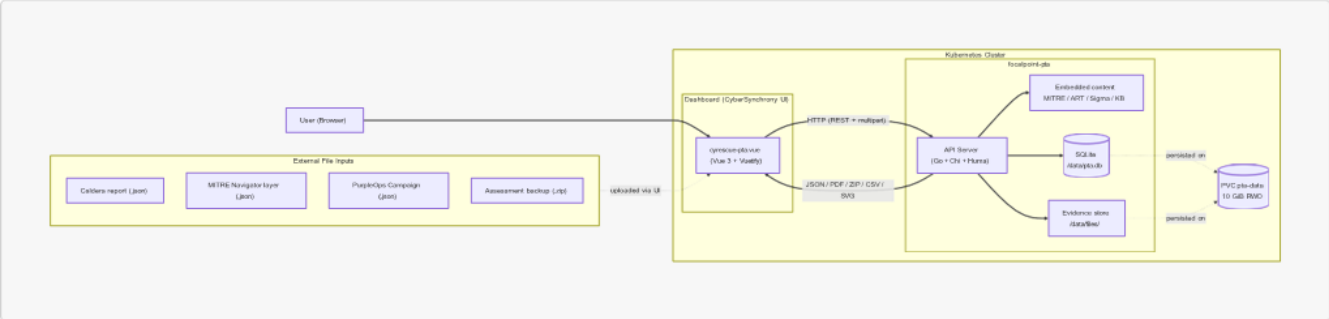

Figure 5: PTA Architecture

5.1 Updates since v1

Since the first version of the Risk Scenario Analysis Toolkit, PTA has evolved from a preliminary description of purple-team assessment functionality into a complete assessment-management service. The current version provides concrete workflows for importing test material, creating and editing assessments, scoring defensive outcomes, attaching evidence, consulting knowledge-base information, and exporting reports.

A major update is the introduction of structured assessment management. Users can create assessments with names and descriptions, restore assessments from ZIP backups, open existing assessments, and delete assessments when needed. This provides persistent management of purple-team exercises and allows the tool to support repeated assessment cycles.

The component has also introduced a rich test-case management workflow. Test cases can be imported from multiple sources, including Caldera reports, MITRE Navigator layers, PurpleOps campaigns, and Atomic Red Team templates. Each test case can be edited through a detailed form that covers MITRE ID, tactic, outcome, state, prevention status, ratings, alert severity, priority, timestamps, objectives, actions, red-team notes, and blue-team notes. The editor sends only the modified fields to the backend, reducing unnecessary updates and improving robustness.

Another important update is the support for evidence management. Each test case can have red-team and blue-team evidence files attached. The system supports multipart uploads, file download, and deletion. This allows the assessment to retain a clear evidence trail that can be used for validation, reporting, and future improvement actions.

PTA also introduces knowledge enrichment for each MITRE technique. The test-case editor can retrieve knowledge-base information and related Sigma rules for the selected MITRE technique. This helps analysts understand the technique, identify relevant detection logic, and connect the exercise outcome with practical defensive guidance.

The reporting and export capabilities have also matured significantly. PTA supports exports in PDF, ZIP, CSV, JSON, and MITRE ATT&CK Navigator layer formats. The PDF report includes a cover, executive summary, per-tactic table, and per-test-case details. The ZIP export provides a full backup including evidence files, allowing assessments to be restored later. CSV and JSON exports support data reuse and analysis, while the Navigator export provides a technique-layer representation using outcome-based scoring.

From a technical perspective, the backend has been implemented as a self-contained Go service using Chi and Huma v2. It uses SQLite as a single-writer persistent store and stores evidence files under a dedicated filesystem path. Database migrations are managed through goose. The backend also embeds the MITRE ATT&CK catalogue, Atomic Red Team templates, Sigma rules, and knowledge-base entries at build time, avoiding runtime dependency on external knowledge sources.

5.2 Integration, interoperability and deployment

PTA is integrated into the CYberSynchrony platform as part of the CYRESCUE module. The frontend is implemented using Vue 3, TypeScript, Vite, and Vuetify 3. The main dashboard page provides the full assessment interface, including the assessments table, template browser, five-tab assessment dialog, nested test-case editor, file upload controls, and export buttons. A typed service library wraps the PTA REST API and supports interaction with the backend. Dialog-scoped state is managed through dedicated composables for assessment details and test-case editing.

The backend is provided by the focalpoint-pta service. It exposes REST and multipart endpoints for assessment creation, assessment import, template import, test-case editing, evidence upload and download, statistics, coverage visualisation, and exports. The API layer includes request handling, correlation and logging middleware, and file-stream support for evidence and exported reports.

The PTA architecture consists of several layers. The dashboard layer manages user interaction, assessment navigation, test-case editing, file upload, and file download. The PTA API layer provides the backend endpoints and OpenAPI documentation in development environments. The embedded content layer provides MITRE ATT&CK techniques, Atomic Red Team templates, Sigma rules, and knowledge-base entries without requiring runtime fetches. The SQLite store maintains durable state for assessments, test cases, multi-entities, tags, and relations. The evidence filesystem stores red-team and blue-team evidence files per test case. Finally, the persistent volume claim provides durable storage that survives pod restarts and recreate rollouts.

Table 6: PTA Layer / Responsibility

Layer	Responsibility
Dashboard (Vue)	Assessments table, templates browser, 5-tab assessment dialog, test-case editor, file upload + download.
PTA API	REST + multipart endpoints, OpenAPI docs (dev only), correlation/log middleware, file-stream evidence.
Embedded content	MITRE techniques, ART templates, Sigma rules, KB entries — compiled into the binary, available offline.

PTA supports several import and export paths, which strengthen its interoperability with existing offensive and defensive tooling. Caldera reports can be imported to create sources, tags, and test cases based on executed abilities. MITRE Navigator layers can be imported to generate pending test cases per technique. PurpleOps campaign files provide a migration path from the legacy stack. ZIP backups allow full restoration of assessments, including test cases, multi-entities, and evidence files. Exports are supported in PDF, ZIP, CSV, JSON, and MITRE Navigator formats.

The component is deployed as a containerised service in Kubernetes. The Docker build uses a multi-stage approach, producing a static Go binary and running it on a distroless runtime image. The service runs as a non-root user in production. Because SQLite is used as a single-writer store, the Kubernetes deployment uses one replica and a Recreate strategy. The service exposes liveness and readiness checks on /api/health and is internally exposed through a ClusterIP service mapping port 80 to container port 8080. Persistent storage is provided through a 10 GiB ReadWriteOnce PersistentVolumeClaim mounted at /data. Configuration is provided through a ConfigMap. The standard resource footprint is 100m CPU and 128Mi memory requested, with limits of 500m CPU and 512Mi memory per pod.

The representative end-to-end flow begins when the user creates a new assessment through the dashboard. The backend stores the assessment in SQLite and returns the assessment identifier. The user can then import a Caldera report, which the backend parses to create sources, tags, and test cases based on the executed abilities. The user opens a test case, updates its state and outcome, records blue-team notes, and optionally attaches evidence files. The user can then switch to the Coverage tab, where the backend groups outcomes by tactic and renders the SVG heatmap. Finally, the user can export the assessment as a PDF report or in another supported format.

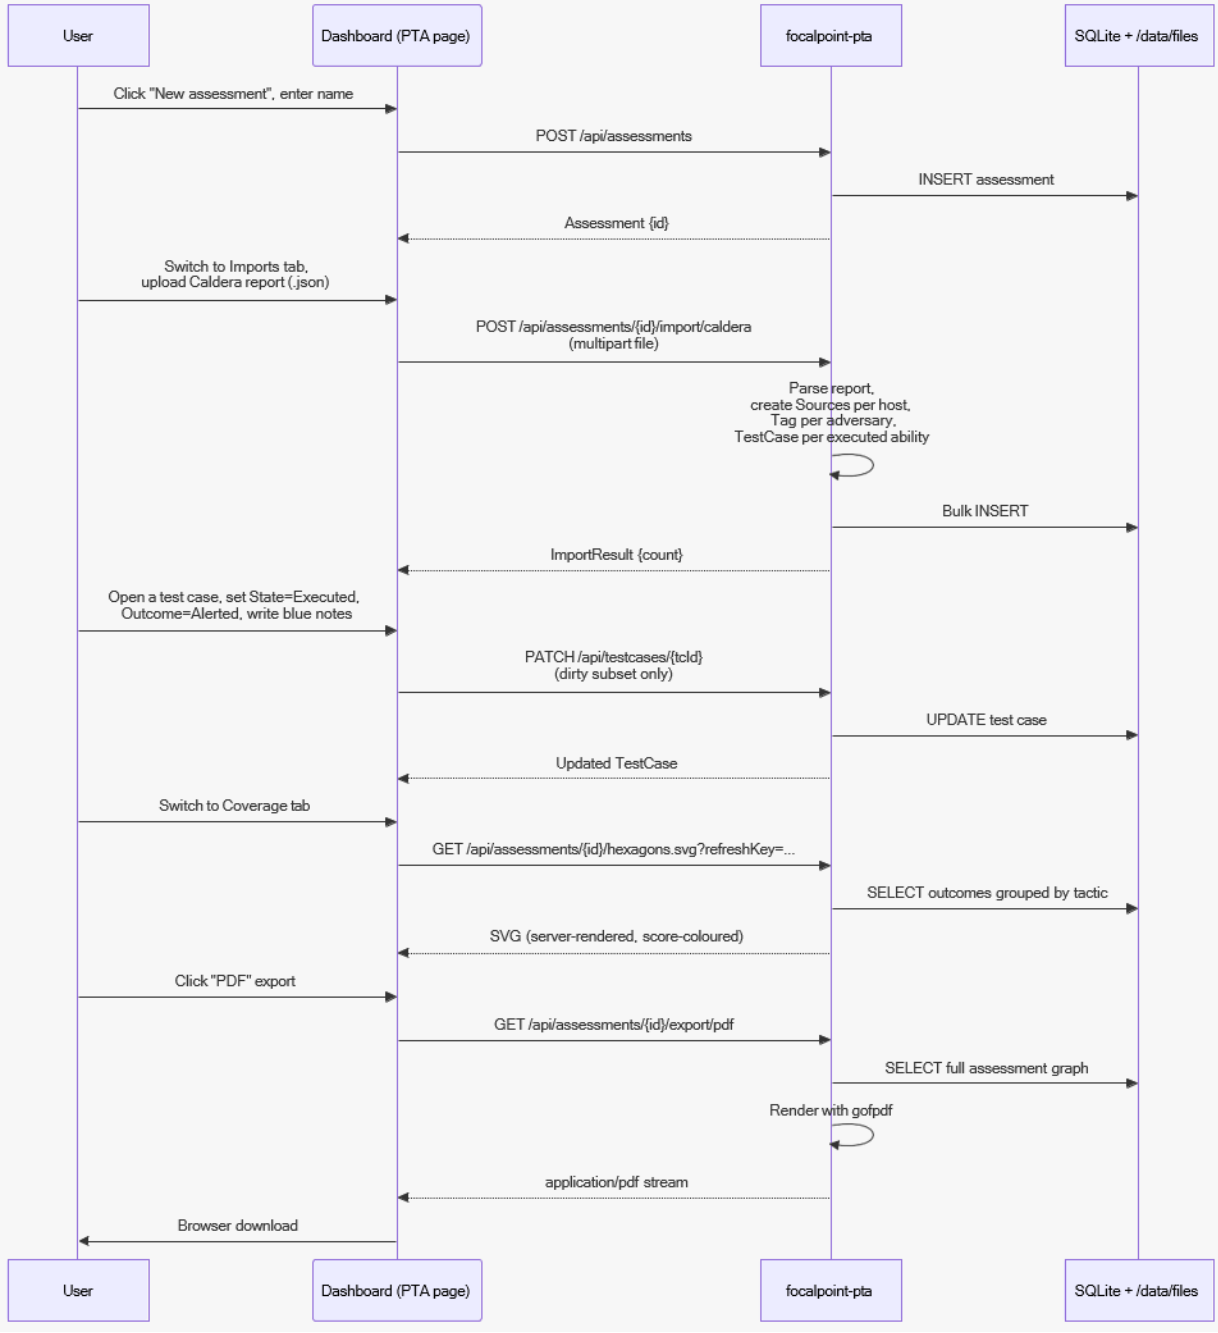

Figure 6: PTA End-to-End Flow

6 KPIs Mapping

This chapter maps monitoring suite components documented in this deliverable to the project KPIs to which they contribute. Measured KPI values are reported through WP8 deliverables.

Pilot operations start from M25, with the first are expected during M25 and M26.

6.1 KPI3.7 Risk Mitigation Index

KPI ID	KPI Name	Owner
3.7	Risk Mitigation Index	MAG
Target		
Achieve 20% reduction in identified risks		
Measurement Methodology		
Components involved	**Participating partners & Role in KPI Measurement**	**Evidence collection process & validation approach**
Risk Assessment Suite (DCO, CSPA, ASA, PTA, SBS	FP operates a set of tools that feed both the initial and residual risk scores. The Attack Surface Assessment (ASA) identifies exposed assets, open services, and CVEs that contribute to initial likelihood. The Cybersecurity Posture Assessment (CSPA) assesses the security posture of the affected asset and informs likelihood once controls are applied. Defensive Cyber Operations (DCO) documents the defensive playbooks whose presence in the residual configuration reduces likelihood. PTA and SBS run purple-team and breach-simulation exercises whose outcomes inform impact scoring before and after mitigation.	Tool reports from Defensive Cyber Operations (DCO), Cybersecurity Posture Assessment (CSPA), AttackSurface Assessment (ASA), Purple team Assessment (PTA), and Security Breach Simulation (SBS); documented likelihood and impact values per scenario at both initial and residual stages; remediation status of each mitigation action; simulation outcome records; analyst notes mapping tool's output to the corresponding scoring change in the register.
trustSense	TRUST uses trustSense to gather participatory security assessment data from the user/operator population affected by each scenario, providing input on perceived risk, behavioural exposure, and observed security gaps. This feeds the likelihood component of the initial score and the residual score where awareness or policy interventions are part of the mitigation.	trustSense survey exports, response statistics, identified behavioural gaps, pre- and post-intervention readings, and analyst notes linking the trustSense output to the likelihood scoring change.

Kitsune Search / Monitoring	MER uses Kitsune Search and Kitsune Monitoring to retrieve external threat intelligence relevant to the assets in scope of each scenario, including data leak exposures, dark-web mentions, and infostealer/botnet activity. Presence and recency of such intelligence inform the initial likelihood score, and reduced exposure after mitigation contributes to the residual likelihood score.	Kitsune Search query exports, Kitsune Monitoring alert logs, lists of exposed credentials or asset mentions, screenshots, and analyst notes correlating each finding with the likelihood scoring change in the register.

Methodology Description

The owner will maintain the Risk Scenario Register (shared with KPI3.6), where each risk scenario receives two scores. The initial risk score is calculated before any mitigation is applied or recommended. The residual risk score is calculated after the mitigation actions described in the row are deployed in the pilot environment. Both scores use the formula Likelihood × Impact, with Likelihood and Impact each scored from 1 to 5 by the responsible partner against a published rubric with explicit anchors for each level. Each scenario row will also record the affected pilot, asset, threat type, list of mitigation actions, deployment status of each action (deployed / recommended / planned), validating pilot owner, and an evidence link to the contributing tool outputs.

The Risk Mitigation Index for a single scenario is calculated as (Initial Risk Score − Residual Risk Score) / Initial Risk Score × 100, expressed as a percentage. The KPI value is the average of the index across all scenarios with a validated residual score. A scenario is considered validated when the responsible pilot owner (AMC, EGL, or UPRC depending on which pilot the scenario belongs to) has signed off that the residual score reflects mitigation actually deployed in the pilot environment, evidenced by the agreed sign-off artefact.

The benchmark of 20% average risk reduction is met when the average index across all validated scenarios is at or above that threshold. Scenarios where the mitigation has been recommended but not deployed are flagged in the register and excluded from the KPI calculation. Where the same scenario is updated over time (additional mitigation deployed in later iterations), the residual score is updated in place and the calculation re-run on the latest values.

KPI Progress Status

Current Value/Target	**Date of Measurement**	**Status**	**Next Actions** (if required)
0% / ≥20% (no validated scenarios yet)	—	On track	Publish the Likelihood × Impact scoring rubric with anchors for each level. Run a calibration round with FP, TRUST, and MER on a sample scenario before full collection begins. Agree the pilot-owner sign-off artefact with AMC, EGL, and UPRC. Instantiate the shared Risk Scenario Register with the D3.4 / KPI3.6 author.

Challenges, Deviations and Corrective Actions

A key challenge for this KPI is the subjectivity of the Likelihood × Impact scoring, since two analysts assessing the same scenario can reach different totals. The deliverable will publish a scoring rubric with explicit anchors for each value from 1 to 5 on both axes, and a short calibration exercise will be run with FP, TRUST, and MER

before the first scenarios are scored, to align the scoring discipline across partners. A related concern is scenario selection bias: the average across "all validated scenarios" can be inflated if partners preferentially include scenarios where mitigation was effective. To address this, the methodology requires that every scenario counted under KPI3.6 must also receive a residual score and be included in the KPI3.7 calculation if its residual score is validated, with no separate selection step. A third open item is the distinction between mitigation that is actually deployed in the pilot environment and mitigation that is only recommended. Only deployed mitigation produces a validated residual score; recommended-but-not-deployed mitigation will be tracked in the register but excluded from the average, so that the reported reduction reflects real-world risk change rather than advisory output. The validation gate itself, defined as the artefact that constitutes pilot owner sign-off, needs to be agreed with AMC, EGL, and UPRC before scoring begins.

7 Conclusions

The Risk Scenario Analysis Toolkit v2 represents the successful maturation of the project's risk-evaluation framework from the foundational concepts established in version 1 to a production-ready suite fully integrated into the CYberSynchrony platform. By operationalizing the synergy between Kitsune Monitoring, CSPA, ASA and PTA, the toolkit provides a multi-layered defense mechanism that moves beyond static alert reporting to evidence-based situational awareness.

The transition to this consolidated release is defined by several technical milestones: Kitsune Monitoring has evolved from a backend API to a web-based visual module for tracking dark web leaks and botnets; the Cybersecurity Posture Assessment (CSPA) now offers automated asset-to-posture mapping against the MITRE D3FEND and ATT&CK domains; the Attack Surface Assessment (ASA) provides agentless vulnerability discovery correlated with real-time NVD and CVE data; and the Purple Team Assessment (PTA) provides a complete management surface for simulating and scoring complex adversary behaviors.

D3.4 establishes also the technical verification necessary to satisfy the project's Key Performance Indicators, specifically contributing to the target of a 20% reduction in identified risks (KPI 3.7) and the analysis of at least 12 customized risk scenarios (KPI 3.6).

As the consortium moves into the M25+ pilot validation phase across the healthcare, municipal, and academic sectors, this toolkit provides the operational agility required to document realistic "what if" situations and maintain a tamper-evident evidence plane for cross-border collaboration.

This release ensures that Member States can share security insights with confidence, supported by a framework that is both forensically valid and technically mature.

8 References

[1] Wazuh – Open Source Security Platform. [Online]. Available: https://wazuh.com/

[2] MITRE D3FEND. [Online]. Available: https://d3fend.mitre.org/

[3] MITRE ATT&CK. [Online]. Available: https://attack.mitre.org/

[4] NVD – National Vulnerability Database. [Online]. Available: https://nvd.nist.gov/

[5] GitLab registry. [Online]. Available: http://registry.gitlab.com

[6] CVE - Common Vulnerabilities and Exposures. [Online]. Available: https://cve.mitre.org/

[7] CWE - Common Weakness Enumeration. [Online]. Available: https://cwe.mitre.org/

[8] MITRE Caldera. [Online]. Available: https://caldera.mitre.org/

Annex I: PoC and Prototypes (Kitsune Monitoring)

D3.4 integrates the specific work performed to achieve integration and interoperability with the CYberSynchrony platform. Deployment of a dedicated web-based graphical user interface (GUI) fully integrated into the project's centralized dashboard, allow SOC analysts to move beyond raw JSON data processing.

Integrated Kitsune Front-End and Intelligence Reach

This paragraph provides a detailed overview of the processes and information accessible through the Kitsune Monitor front-end module integrated within the CYberSynchrony dashboard.

- **Threat Landscape Overview (Homepage)**

The integrated homepage serves as the primary entry point for global and localized situational awareness.

It provides real-time statistics on the current global threat volume, including counts for tracked Threat Actors, Malware samples, and Indicators of Compromise.

The dashboard prioritizes the most active threats and campaigns, allowing analysts to immediately see which actors (e.g., Turla, APT28) are most active and which sectors (e.g., Government, Technology) are currently under heavy attack

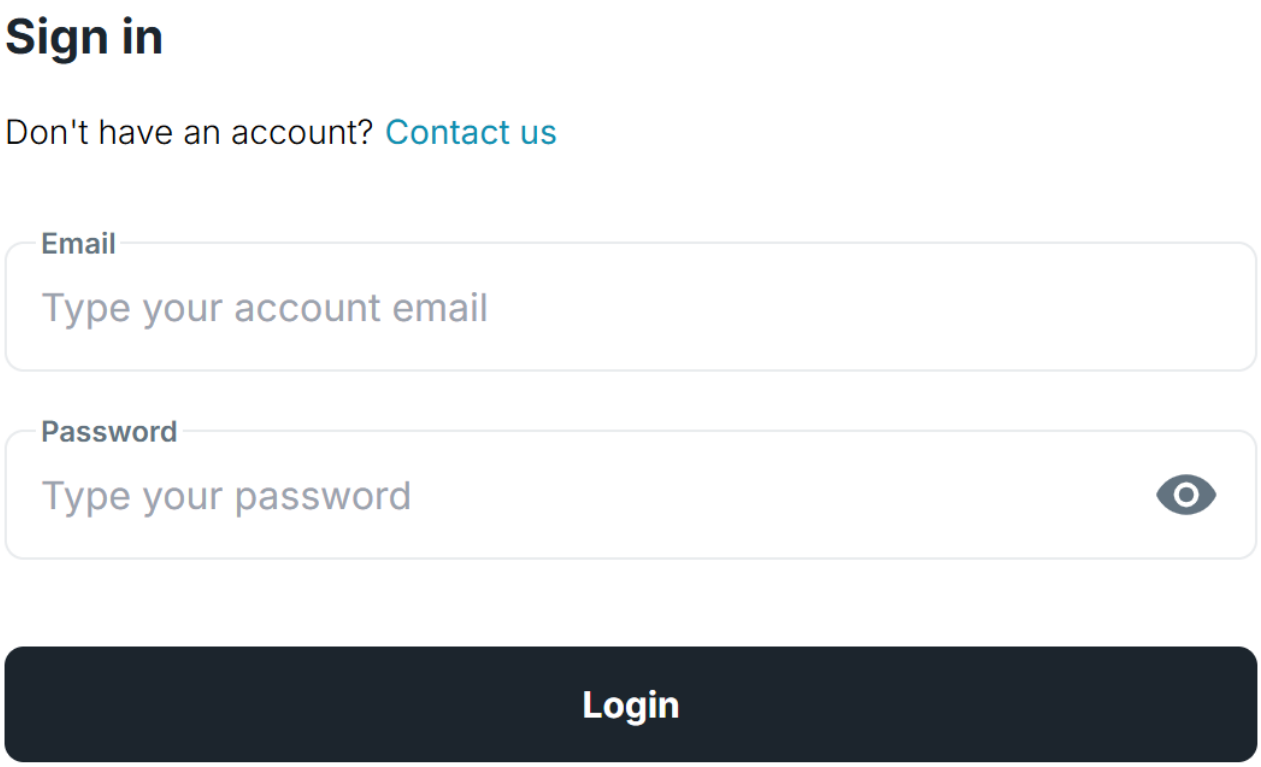

Figure 7: Kitsune Login

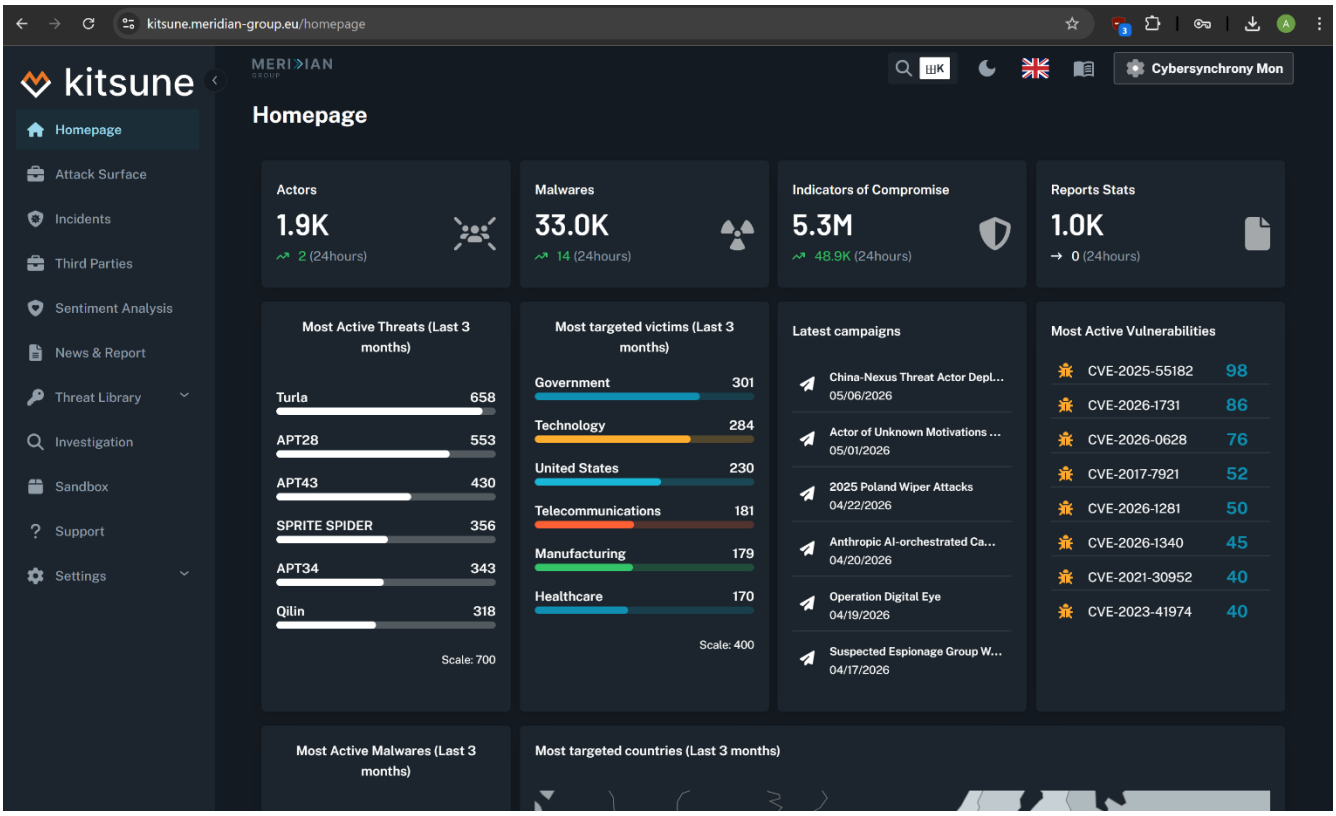

Figure 8: Kitsune Homepage

- **Investigation and Incident Triage**

The integrated Incidents module allows analysts to manage findings from the Kitsune intelligence engine as formal security cases:

- **Information Leakage Views**: detailed reporting on leaked employee credentials found on public paste sites or underground forums, including plain-text exposure fragments and remediation recommendations like password rotation.
- **Exposure Analysis**: tracking of outdated web services exposed on public IPs or phishing domains spoofing the corporate brand.
- **Severity Classification**: every incident is assigned a severity level (Critical, High, Medium, Low) to facilitate rapid response prioritization.

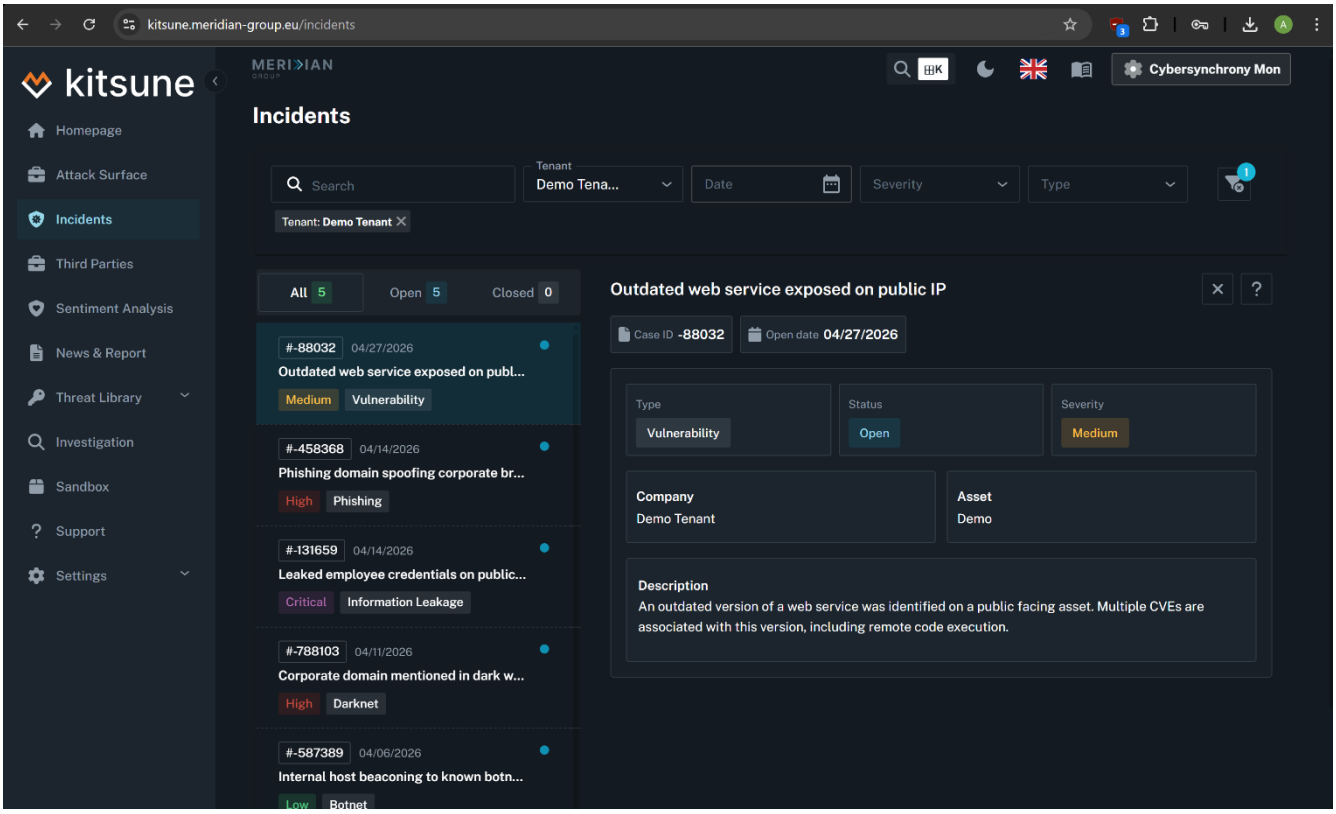

Figure 9: Kitsune Incidents Overview

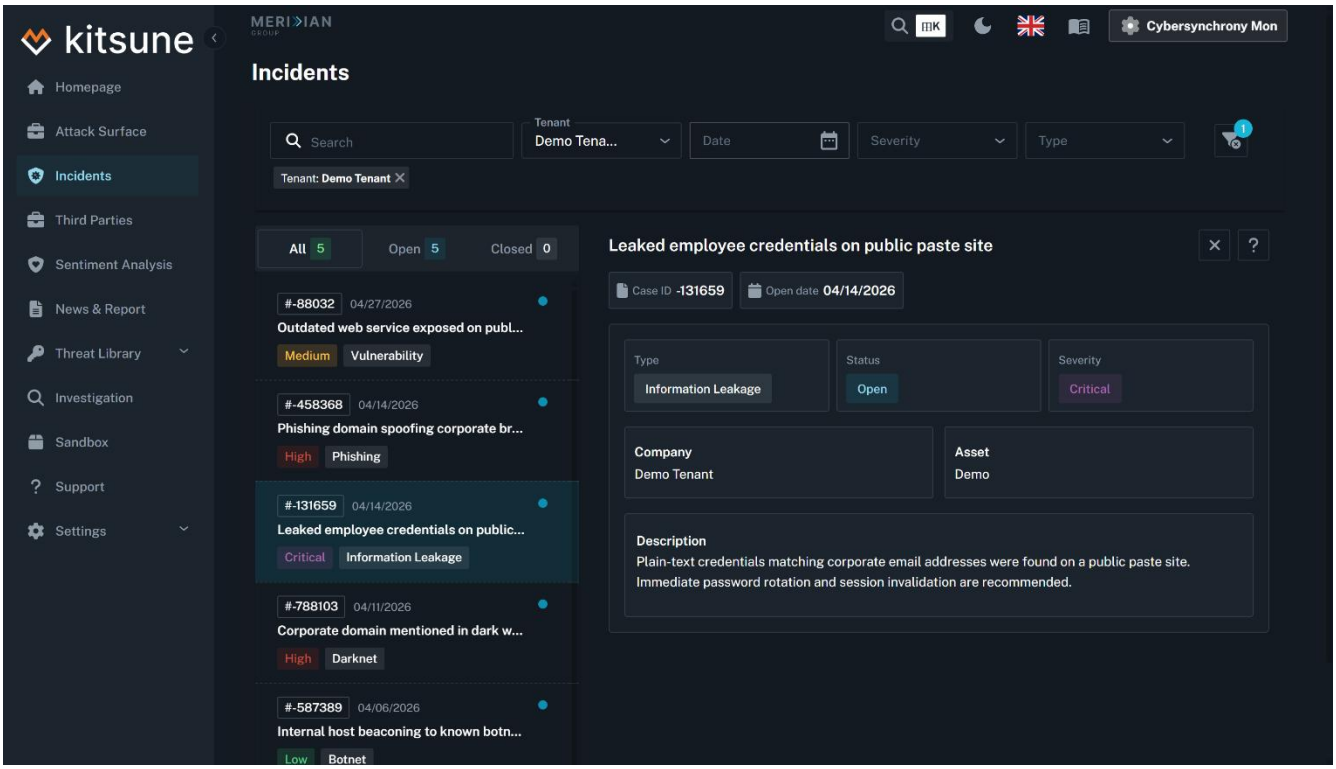

Figure 10: Kitsune Incidents Details

Annex II: Integration of ASA, CSPA into the CYberSynchrony Platform

This section provides a visual walkthrough of the matured and operationally integrated dashboards for the Attack Surface Assessment and Cybersecurity Posture Assessment tools.

These images illustrate the user-facing interface for the vulnerability-discovery capability of the CYBERGOPLUS module.

They showcase the high-level assessment summary, presenting the total count of discovered hardware, applications, and services, with a granular "Detailed" view, where analysts can review individual resolved artifacts, their CPE strings, and a scrollable grid of associated CVEs enriched with severity badges and CVSS base scores for rapid triage:

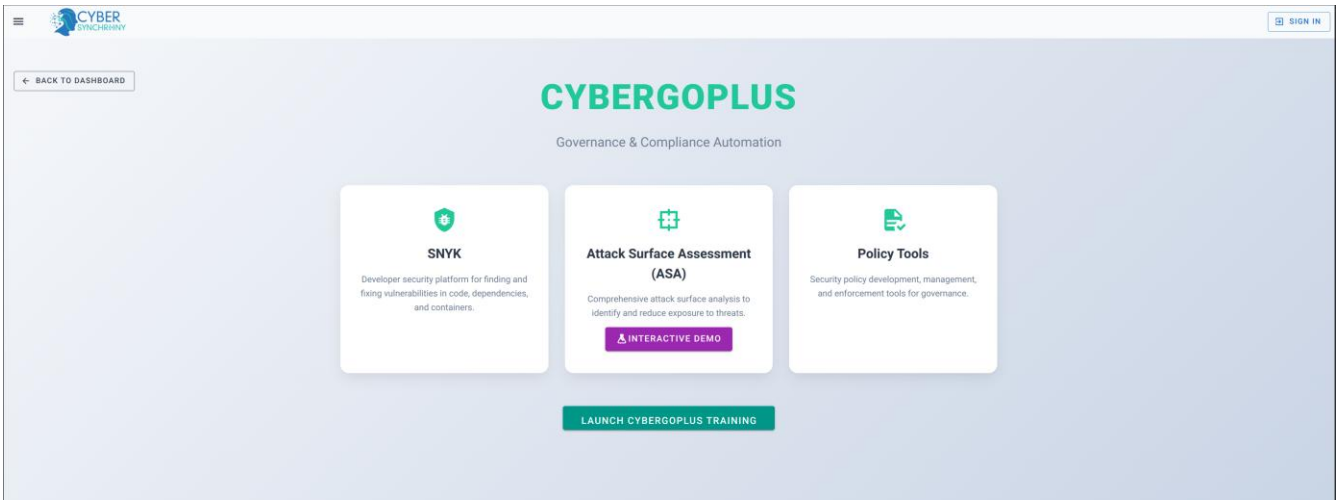

Figure 11: ASA login access

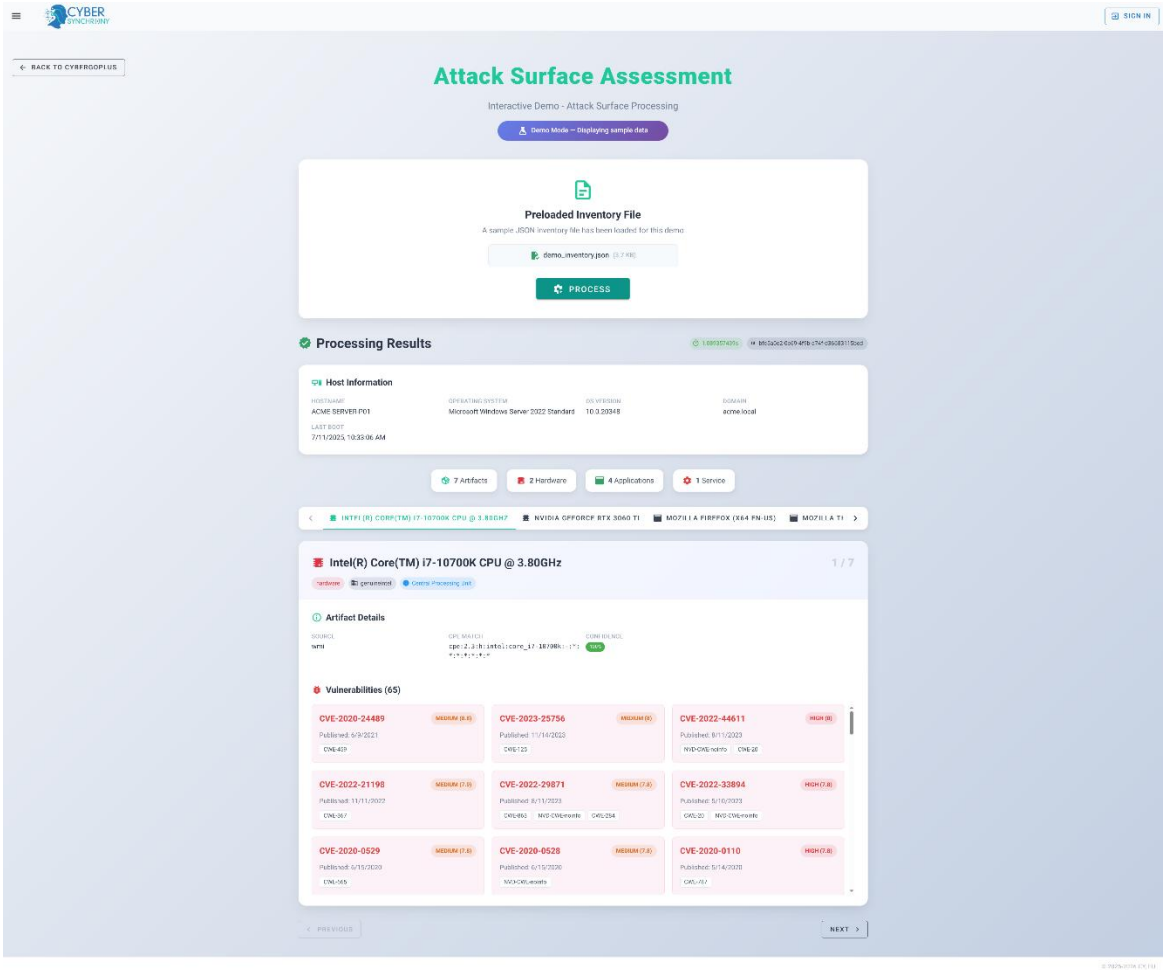

Figure 12 : ASA Detailed View: Operational Posture Mapping

Following figures depict the operational posture-mapping interface within the CYRESCUE module.

They demonstrate the detailed analysis layer, utilizing nested expandable panels to show how each identified asset correlates with MITRE D3FEND nodes (detection strategies and mitigations) and the MITRE ATT&CK techniques that may target them.

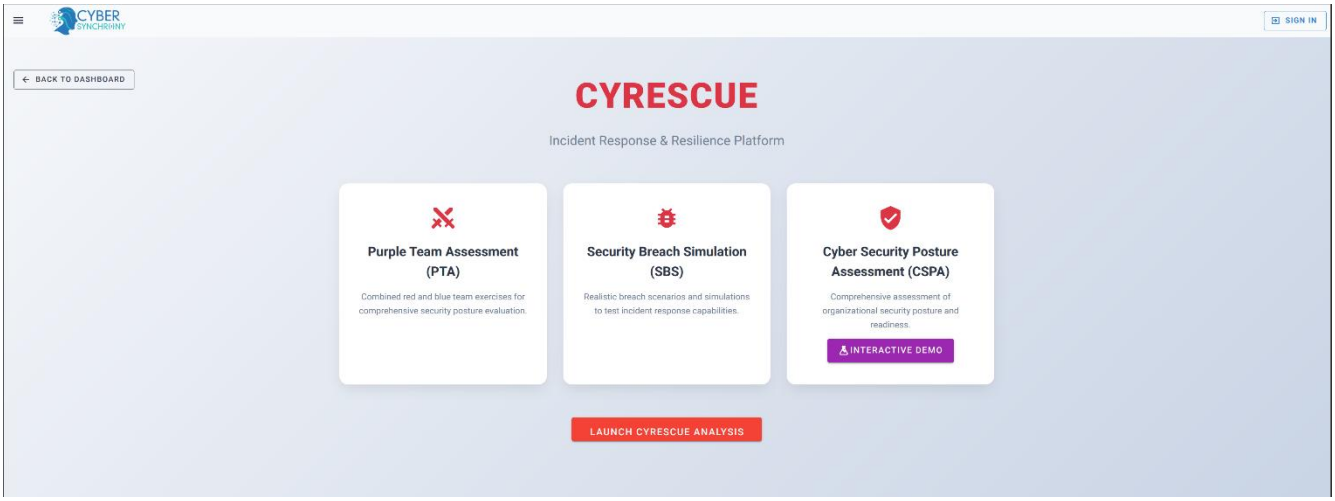

Figure 13: CSPA login access

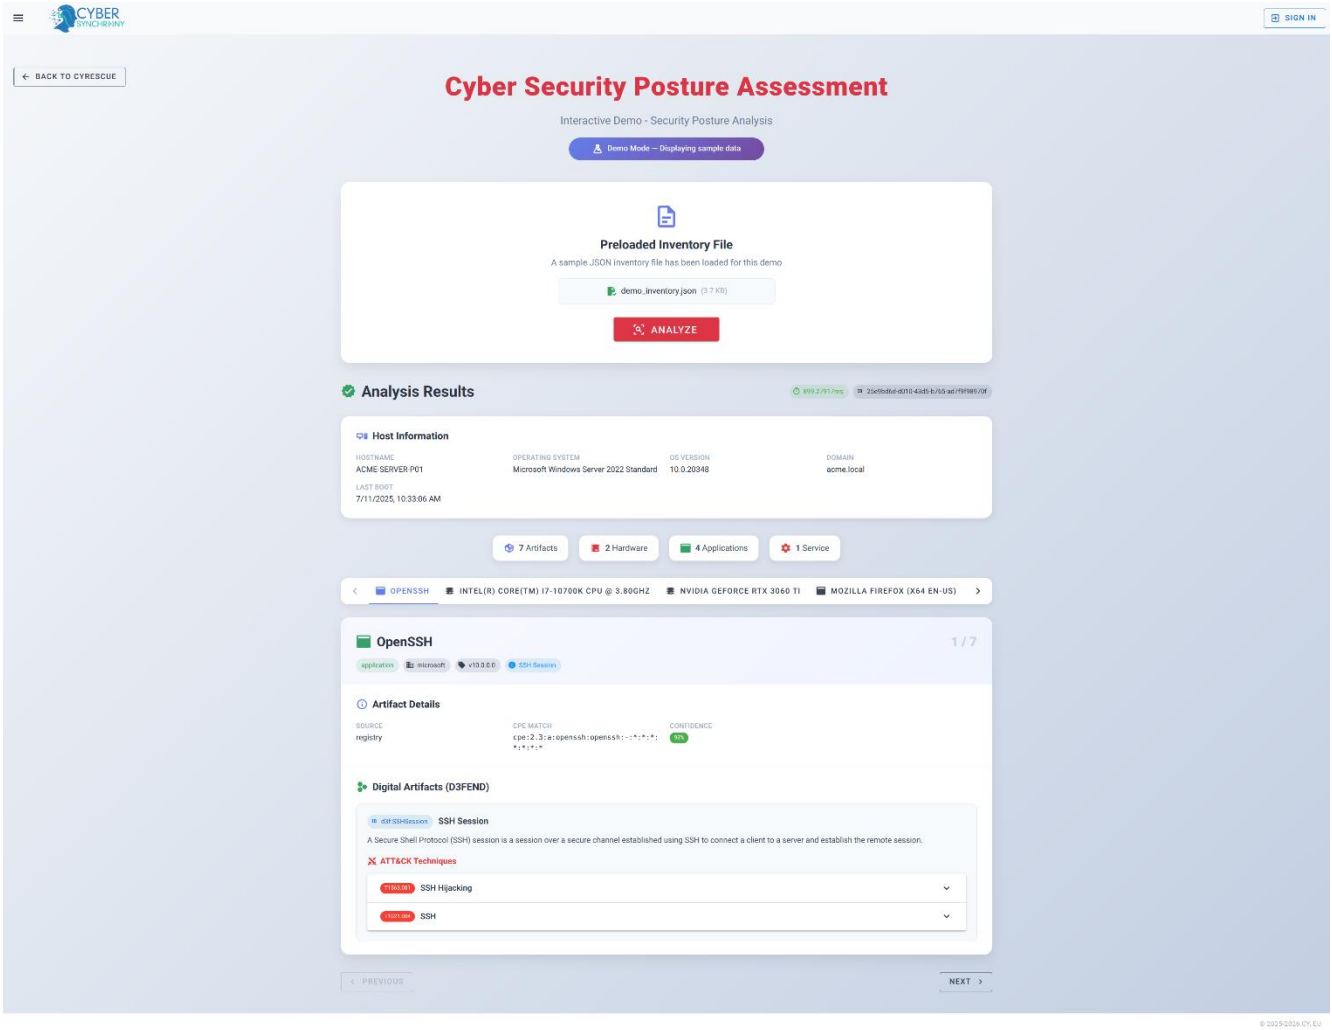

Figure 14: CSPA Detailed View: Operational Posture Mapping