Harmonizing People, Processes, and Technology for Robust Cybersecurity

D6.8 Compliance Monitoring 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	D6.8 Compliance Monitoring		
Work Package	WP6 – CYBERWISE-CYBERGOPLUS		
Contractual due date	31/08/2026	**Actual submission date**	24/07/2026
Type	Report	**Dissemination Level**	PU
Lead Beneficiary	IANUS		
Responsible Author	Demetris Antoniou (IANUS)		
Contributions from	BYTE, MAG, DSA, eBOS		

Funded by the European Union

Funded by the European Union under Grant Agreement no. 101158555. 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	15.04.2026	Initial deliverable structure (evolution of D6.4 v1)	Demetris Antoniou (IANUS)
0.2	27.04.2026	Draft text version	Demetris Antoniou (IANUS)
0.3	11.05.2026	Contribution on Real-Time compliance oversight (Task 6.7)	George Karavokiros (BYTE)
0.4	18.05.2026	Contribution on integration logic and interfaces	Kostantinos Andrikos (MAG)
0.5	29.05.2026	Update and implementation of DSA requirements	George Achilleos (DSA)
0.6	22.06.2026	90% of the deliverable ready for peer-review, obtained input from participants	Demetris Antoniou (IANUS)
0.7	17.07.2026	Peer-Review	KETYAK, UPRC
0.8	20.07.2026	Addressed peer-review comments & submitted for QA	Demetris Antoniou (IANUS)
0.9	23.07.2026	QA review	Eleni Keliri (eBOS)
1.0	30.07.2026	Deliverable ready 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

© CYberSynchorny, 2025

Table of Contents

1 Executive Summary 10
2 Introduction 12
2.1 Purpose of the Deliverable 12
2.2 Mapping CYberSynchrony Outputs 12
2.3 Linkage to other project outputs 13
2.4 Deliverable Structure 14
3 Overview of WP6 and Task 6.7 16
3.1 WP6 Scope and Cross-Sector Orientation 16
3.2 Real-Time Compliance Oversight 16
3.3 Dependencies with Tasks 6.1–6.6 17
3.4 Deliverable Evolution and Key Outputs 18
3.4.1 Foundations Carried Forward from v1 (Definition and Architecture) 18
3.4.2 Advances Introduced in v2 (Specification, Integration and Validation Framework) 18
4 Compliance Framework Foundations 20
4.1.1 NIS2 Directive and National Implementation Frameworks 20
4.2 Sector-Specific Requirements (Healthcare, Governance, Academia) 22
4.2.1 Healthcare Sector Requirements 22
4.2.2 Public Administration Sector Requirements 24
4.2.3 Academia Sector Requirements 25
4.3 Harmonised Cross-Sector Principles 26
4.3.1 Principle 1: Risk-Driven Governance 27
4.3.2 Principle 2: Asset-Centric Protection 27
4.3.3 Principle 3: Continuous Threat Detection 28
4.3.4 Principle 4: Resilience and Recovery 28
4.4 Mapping of Principles to CYberSynchrony Responsibilities 29
5 CYberSynchrony Compliance Monitoring Approach (High-Level) 32
5.1 Objectives of Compliance Monitoring 32
5.2 Compliance Domains and Assessment Areas 32
5.3 Indicators, Metrics and Maturity Levels 33
5.3.1 Key Compliance Indicators (KCIs) 33
5.3.2 Mapping Controls to KPIs 34
5.3.3 Maturity Levels 36
5.3.4 KPI Measurement Methodology 36
5.4 Evidence Bases: Policy, Implementation, and Operational Artefacts 37
5.4.1 Policy and Governance Evidence 38
5.4.2 Implementation and Configuration Evidence 38
5.4.3 Operational and Runtime Evidence 39
6 System Specifications and Integration Requirements 42
6.1 Functional Requirements 42
6.1.1 Policy and Governance Management Functions 42
6.1.2 Automated Auditing and Assessment Functions 42
6.1.3 Findings Management and Reporting Functions 43
6.2 Technical Specifications 43
6.2.1 Architecture Pattern 43
6.2.2 Technology Stack 43
6.2.3 Data Storage Strategy 44
6.2.4 Security Standards 44
6.3 Integration Specifications 44

© CYberSynchorny, 2025

6.3.1 Integration with Evidence Sources 44
6.3.2 Integration with Identity Services 46
6.3.3 Integration with Presentation Layer 47
7 End-to-End Compliance Workflows 49
7.1 The Compliance Monitoring Loop 49
7.1.1 The Loop in API Terms 50
7.2 Phishing Incident Detection and Response 51
7.3 Vulnerability Detection and Remediation 52
7.4 Identity and Access Compliance Monitoring 53
7.5 Governance and Risk Verification 54
7.6 Resilience and Integrity Verification 55
7.7 Control Coverage Summary 56
8 Per-Pilot Integrated Narratives 58
8.1 UC1: Healthcare (Apostolos Loucas Medical Center) 58
8.2 UC2: Public Administration (EGL) 59
8.3 UC3: Academia (UPRC) 61
8.4 Cross-Pilot Observations 62
9 Ethics, Privacy and Data Protection 64
9.1 Personal Data Processed and Lawful Basis 64
9.2 Data Minimization and Protective Measures 65
9.3 Privacy by Design and by Default 66
9.4 Ethics and Proportionality 66
10 Conclusions 68
10.1 Summary of Contributions 68
10.2 Limitations 69
10.3 Risk Reflection 69
10.4 Outlook and Path to Validation 70
11 References 71
Annex I: Technical Interface Specifications 72
Annex II: Standard Adoption Baseline Evidence (KPI 9.2) 85

© CYberSynchorny, 2025

List of Figures

Figure 1: Harmonised Cross-Sector Compliance Principles mapping regulatory concepts to CYberSynchrony platform responsibilities..... 27

Figure 2: Functional Compliance Monitoring Architecture detailing the flow from raw telemetry to stakeholder reporting..... 41

Figure 3: CYBERGOPLUS Technical Architecture. This diagram illustrates the module's microservices deployment within Kubernetes, detailing data ingestion from external sources (CYBRITE, CROSS-CORE, MISP, and Kitsune), internal processing via the Compliance Engine..... 48

Figure 4: Phishing Incident Response Playbook. 51

Figure 5: EU enterprise cybersecurity control adoption by size, 2024 (Eurostat) 86

Figure 6: NIS2 implementation milestones, 2018 to 2026..... 87

Figure 7: Member States under Commission action for incomplete NIS2 transposition 87

List of Tables

Table 1: Adherence to CYberSynchrony's GA..... 12

Table 2: Links to other Project Outputs/Work 13

Table 3: Mapping of Platform Functionality to Regulatory Evidence 20

Table 4: Healthcare-Specific Compliance Indicators 23

Table 5: Public Administration-Specific Compliance Indicators..... 24

Table 6: Academia-Specific Compliance Indicators..... 26

Table 7: Mapping of Harmonised Principles to Platform Responsibilities 29

Table 8: DSA Derived Controls and Description 30

Table 9: Operational Compliance Indicators by Evidence Source..... 33

Table 10: Comprehensive Compliance Indicator Map 35

Table 11: Personal Data Categories Processed by CYBERGOPLUS..... 64

Table 12: Comparative Standard Adoption by Critical Sector..... 85

© CYberSynchorny, 2025

Acronyms and abbreviations

Abbreviation / Term	Description
AM	Asset Management
API	Application Programming Interface
BFF	Backend for Frontend
BIA	Business Impact Analysis
BYOD	Bring Your Own Device
CAPA	Corrective and Preventive Action
CISO	Chief Information Security Officer
CNI	Container Network Interface
CVE	Common Vulnerabilities and Exposures
CYBERRA	Cyber Threat Intelligence and Customized Risk Analysis
CYBRITE	Proactive Cyber Risk Monitoring and Incident Response
CROSS-CORE	Cross-Border Cyber Incident Cooperation & Response
CYBERGOPLUS	Cybersecurity Governance and Compliance
CYBERWISE	Cybersecurity Consciousness: Empowering Society Through Awareness
CYRESCUE	Cybersecurity Resilience and Incident Response Enhancement
DCO	Defensive Cyber Operations
DDoS	Distributed Denial of Service
DPO	Data Protection Officer
DSA	Digital Security Authority (Cyprus)
EDR	Endpoint Detection and Response
EHR	Electronic Health Records
EIM	Event and Incident Management
FIM	File Integrity Monitoring
GDPR	General Data Protection Regulation
GOV	Governance
GRC	Governance, Risk, and Compliance
gRPC	gRPC Remote Procedure Calls
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure

© CYberSynchorny, 2025

IaC	Infrastructure as Code
IAM	Identity and Access Management
IDS	Intrusion Detection System
IoMT	Internet of Medical Things
IP	Internet Protocol
IPS	Intrusion Prevention System
ISO	International Organization for Standardization
JSON	JavaScript Object Notation
JWT	JSON Web Token
K8s	Kubernetes
KCI	Key Compliance Indicator
KPI	Key Performance Indicator
LMS	Learning Management System
MFA	Multi-Factor Authentication
MTTR	Mean Time to Remediate / Respond
NIS2	Network and Information Security Directive 2
NIST	National Institute of Standards and Technology
NS	Network Security
OIDC	OpenID Connect
OS	Operating System
PHI	Protected Health Information
PII	Personally Identifiable Information
RBAC	Role-Based Access Control
REST	Representational State Transfer
RM	Risk Management
SAST	Static Application Security Testing
SIEM	Security Information and Event Management
SLA	Service Level Agreement
SNYK	Security tool for vulnerability scanning
SOC	Security Operations Center
SS	System Security
SSO	Single Sign-On

© CYberSynchorny, 2025

SQL	Structured Query Language
TA	Training and Awareness
TLS	Transport Layer Security
TLP	Traffic Light Protocol
UI	User Interface
URL	Uniform Resource Locator
VM	Vulnerability Management
WP	Work Package
XML	Extensible Markup Language

© CYberSynchorny, 2025

1 Executive Summary

Task 6.7 “Real-Time Compliance Oversight” plays a crucial role within CYberSynchrony by establishing the mechanisms, indicators and monitoring pipelines required to ensure that organisations adopting the CYberSynchrony ecosystem can demonstrate continuous alignment with cybersecurity and data protection regulations. As regulatory frameworks such as NIS2 and GDPR, and technical standards such as NIST 800 and ISO 27001, shift towards evidence-driven compliance, Task 6.7 introduces a structured methodology for real-time compliance oversight. As organisations in Healthcare, Public Administration and Academia digitise their operations, they become subject to an increasingly complex set of regulatory requirements, and a persistent challenge remains the gap between written policy and technical reality: organisations often hold comprehensive security policies on paper yet lack the visibility to verify, in real time, that those policies are enforced across their infrastructure.

Deliverable 6.8 Compliance Monitoring v2 is the second and final version of the Compliance Monitoring report, developed under Work Package 6. It builds directly on the framework established in D6.4 (v1) and advances it from the definition and architecture stage to a fully specified operational design. The deliverable retains the v1 objective of the CYBERGOPLUS module, the dedicated component that continuously monitors, audits and visualises an organisation's compliance posture, and extends it with the concrete specification, additional evidence sources and measurement methodology required to operationalise it.

Version 1 served three interconnected purposes: defining the "To-Be" technical state by translating the 15 static Policy Tools of D6.3 into measurable technical requirements; establishing the "As-Is" baseline of the pilot use cases from the D2.2 analysis and the M18 surveys; and architecting the CYBERGOPLUS solution around open-source components such as Wazuh and Keycloak. Version 2 preserves these foundations and advances them in the following ways:

1. **Specifying the Monitoring Engine**: v2 specifies CYBERGOPLUS at the level of the actual evidence artifacts produced by the integrated components, defining precisely which output of each component satisfies each control retained from v1.
2. **Extending the Evidence Base**: v2 extends the v1 evidence base with additional sources mapped to existing controls, including the threat-intelligence and secure-sharing outputs of MISP and the team trustworthiness assessment produced by the trustSense tool, without introducing new compliance controls.
3. **Defining the Measurement Methodology**: v2 establishes a structured KPI methodology in which each indicator carries an explicit baseline, target population, measurement window, evidence source and associated confidence, and distinguishes preparatory indicators from outcome indicators.
4. **Operationalising the Framework**: v2 advances the component-level descriptions of v1 into end-to-end cross-module workflows, consolidates the sector requirements into integrated per-pilot narratives, and adds a dedicated treatment of ethics, privacy and data protection.

The baseline established in v1 remains the reference point against which the monitoring framework is calibrated. That assessment found high policy adoption across pilot organisations (85% had formally adopted policies), low technical enforcement (35% had automated tools to verify enforcement on endpoints), and a significant reporting burden (71% of D2.2 respondents identified compliance reporting as a critical need). These figures define the starting conditions the v2 design is built to address.

This version completes the design of the compliance monitoring system and defines the methodology and validation framework for its operational assessment within the pilot environments. In line with the project timeline, longitudinal operational validation across the pilots continues beyond this submission. As the final version of the deliverable, due in Month 36, D6.8 confirms that real-time compliance oversight is not a technical luxury but a regulatory necessity under NIS2 and provides the fully specified basis on which that oversight is delivered.

© CYberSynchorny, 2025

A defining characteristic of this version is its progression from a blueprint-level description to an operational specification. Where the first version established the framework's structure and architecture, D6.8 specifies how compliance is produced in practice. Chapter 7 sets out the end-to-end workflows that carry evidence from the integrated components through the CYBERGOPLUS engine to compliance findings and remediation; Chapter 6 and Annex I specify the integration interfaces to the level of concrete endpoints, authentication, and data models; and Chapter 8 applies the framework to the three sectoral pilots as integrated operational narratives. Together, these establish the framework as an operational, evidence-driven system rather than a conceptual design.

© CYberSynchorny, 2025

2 Introduction

2.1 Purpose of the Deliverable

In the rapidly evolving landscape of European cybersecurity, regulatory frameworks such as the **NIS2 Directive** [1] and **GDPR** [2] have shifted the focus from voluntary best practices to mandatory compliance. Organizations operating critical infrastructure or providing essential services must now demonstrate not only that they have security policies in place, but that these policies are actively enforced and effective.

The primary purpose of **Deliverable 6.8 Compliance Monitoring v2** is to advance the operational framework for Task 6.7 (Real-time Compliance Oversight) from definition to a fully specified design. Where v1 established the framework that bridges the gap between governance documentation and technical reality, v2 specifies in concrete terms how that bridge is built, defining how the CYBERGOPLUS module consumes the outputs of the integrated components to verify an organisation's security posture continuously.

Building on v1, which served as the foundational blueprint for the CYBERGOPLUS compliance module, this deliverable retains and develops the three elements that the framework defined:

- **The "What":** The specific regulatory requirements and internal policies (from D6.3) that must be monitored, retained without change from v1.
- **The "How":** The technical architecture and indicators (KCIs) required to translate static policy documents into dynamic, real-time alerts, now specified against the actual evidence artifacts produced by each integrated component.
- **The "Current State":** The baseline assessment of the pilot users established in v1, which remains the reference point against which the monitoring framework and its measurement methodology are calibrated.

Beyond these retained elements, v2 extends the framework with additional evidence sources mapped to the existing controls, a structured measurement methodology, end-to-end cross-module workflows, consolidated per-pilot narratives, and a dedicated treatment of ethics and data protection. None of these additions introduces a new compliance requirement; each develops the framework defined in v1.

2.2 Mapping CYberSynchrony Outputs

This deliverable is a key output of **Work Package 6 (WP6)**, which focuses on "Cybersecurity Consciousness: Empowering Society Through Awareness." Specifically, it addresses the requirements of **Task 6.7**, led by **BYTE**, which mandates the development of "tools that monitor and audit in real-time" to ensure industry-wide compliance. As the second and final version of the Compliance Monitoring report, D6.8 completes the Task 6.7 contribution by advancing the v1 framework to a fully specified operational design. In addition, KPI 9.2. methodology will be developed in this deliverable, for final measurement in in WP8, from the pilot outcomes.

The following table maps the specific Grant Agreement (GA) commitments addressed in this document:

Table 1: Adherence to CYberSynchrony's GA

Component no	Component title	Respective Document Chapter(s)	Justification
DELIVERABLE			
D6.8	**Compliance Monitoring v2**	All	Specifies the operational design, measurement methodology and extended evidence base for compliance

© CYberSynchorny, 2025

			oversight, building on the methodology and findings of v1.
TASKS			
Task 6.7	**Real-time Compliance Oversight**	Chapters 3 to 8	Describes the technical framework for real-time auditing tools, from the regulatory foundations through to the end-to-end workflows and per-pilot application.
KPI			
KPI9.2	**Industry-specific Standard Adoption**	Chapter 5	Provides the measurement methodology and baseline for the industry-specific standard adoption KPI (Objective 9, Task 6.6, IANUS); the measured result is reported in D8.3.

2.3 Linkage to other project outputs

The Compliance Monitoring framework leverages and integrates results from multiple Work Packages across the project. This deliverable is related to other project outputs as it translates the requirements, policies, and technical capabilities defined elsewhere into a structured oversight mechanism. In addition, the contribution of this deliverable and its outcomes to the KPI 9.2. is presented.

Table 2 outlines the specific deliverables and tasks that contribute to or interface with the compliance monitoring activities described in D6.4.

Table 2: Links to other Project Outputs/Work

WP Number	Task Number	Deliverable Number related	Content
WP2	T2.2, T2.4	D2.2 [7]	Provides the cybersecurity landscape analysis, end-user requirements (including reporting needs), and the overarching framework design used to define the compliance baseline.
WP3	T3.2	D3.2[8]	Provides the Risk Scenario Analysis Toolkit and Risk Register structure used for the "Risk-Driven Governance" principle.

© CYberSynchorny, 2025

WP4	T4.1	D4.1[9]	Provides the Integrated Cybersecurity Monitoring Suite (Wazuh, Suricata, Sysmon) which generates the raw technical evidence for asset, network, and vulnerability compliance.
WP4	T4.3	D4.2[10]	Provides the Automated Response Toolkit (DCO Playbooks), used to verify incident response and resilience capabilities.
WP5	T5.2	D5.2[11]	Provides the Evidence Integrity System (Hyperledger), used to verify the immutability of logs and audit trails.
WP6	T6.1	D6.1[12] & D6.5 [20]	Provides the Cyber Awareness Platform and Keycloak Identity Manager, used as the source of truth for Role-Based Access Control (RBAC) and MFA verification.
WP6	T6.4	D6.2[13] & D6.6 [21]	Provides the Cybergate gamified training environment, generating the user performance metrics required to verify staff awareness training compliance.
WP6	T6.5	D6.3[14] & D6.7 [22]	Provides the Cybersecurity Policy Toolkit, containing the 15 static policy templates (e.g., Access Control, Asset Management) that this monitoring module digitizes.
WP8	T8.1	D8.1[15]	Provides the Integration Blueprint and Kubernetes architecture standards that dictate how the compliance module is deployed and connected to the ecosystem.
WP5	T5.1	D5.3 [16]	Provides the Secure Data Exchange Platform (v2), including its MISP-based threat-intelligence sharing and signed, encrypted information exchange, contributing threat indicators and secure-sharing evidence to the monitoring framework.
WP5	T5.2	D5.4 [17]	Provides the Evidence Integrity System (v2), the evolution of D5.2, used to verify the immutability of logs and audit trails.

2.4 Deliverable Structure

This document is organized into the following chapters:

- **Chapter 1: Executive Summary:** Provides an overview of the document, summarising the context, objectives and main findings regarding compliance monitoring within the CYberSynchrony project.
- **Chapter 2: Introduction:** Outlines the purpose of the deliverable, maps the content to project outputs and the Grant Agreement, and details the structure of the document.
- **Chapter 3: Overview of WP6 and Task 6.7:** Describes the scope of Work Package 6, focusing on the role of Task 6.7 in real-time compliance oversight, its dependencies with other tasks, and the evolution of the deliverable from v1 to v2.
- **Chapter 4: Compliance Framework Foundations:** Establishes the regulatory baselines (including NIS2, GDPR and ISO 27001), the sector-specific requirements for Healthcare, Public Administration and Academia, and the harmonised principles underpinning the framework.
- **Chapter 5: CYberSynchrony Compliance Monitoring Approach:** Details the monitoring methodology, covering objectives, compliance domains, indicators, metrics, the KPI measurement methodology, and

© CYberSynchorny, 2025

the evidence bases derived from the policy, implementation and operational layers, including the additional evidence sources integrated in v2.

- **Chapter 6: System Specifications and Integration Requirements:** Defines the functional capabilities, technical architecture, and integration interfaces required for the CYBERGOPLUS module, ensuring alignment with the wider CYberSynchrony ecosystem.
- **Chapter 7: End-to-End Compliance Workflows:** Presents concrete cross-module workflows that trace compliance evidence from its generation in an integrated component, through validation in the CYBERGOPLUS module, to the resulting compliance finding and remediation action.
- **Chapter 8: Per-Pilot Integrated Narratives:** Consolidates the application of the framework to each pilot (Healthcare, Public Administration and Academia) into a single integrated narrative per use case, covering configuration, the workflows exercised, the resulting compliance posture and the operational findings.
- **Chapter 9: Ethics, Privacy and Data Protection:** Assesses the privacy and data-protection implications of the evidence collected for compliance monitoring, including the data-protection impact considerations applicable to the monitoring data.
- **Chapter 10: Conclusions:** Summarises the status of the final release, reflects on the limitations and residual risks of the framework, and outlines the remaining path to operational validation within the pilots.
- **Chapter 11: References**: Lists the bibliographic references used throughout the document.
- **Annex I: Technical Interface Specifications:** Provides the specific data models, JSON schemas, and API payload definitions required for system integration.
- **Annex II: Standard Adoption Baseline Evidence (KPI 9.2):** An analysis of current achievement of adoption of best practices and standards for NIS2 regulated entities.

© CYberSynchorny, 2025

3 Overview of WP6 and Task 6.7

3.1 WP6 Scope and Cross-Sector Orientation

Work Package 6 (WP6), titled "Cybersecurity Consciousness: Empowering Society Through Awareness (CYBERWISE) and Cybersecurity Standardization and Compliance Reinforcement (CYBERGOPLUS)," is designed to address the human and governance layers of the cybersecurity ecosystem. The scope of this work package extends beyond technical implementation to encompass the behavioral and organizational changes necessary to maintain a secure posture.

The primary objectives of WP6 are to:

- Engage diverse societal segments with cybersecurity awareness.
- Utilize gamification techniques to make learning impactful.
- Measure the effectiveness of awareness campaigns using data analytics.
- Cultivate a culture prioritizing secure online behavior.
- Support comprehensive cybersecurity policy implementation.
- Establish unified standards across various sectors.
- Monitor and audit adherence to established cybersecurity standards.
- Ensure industry-wide compliance with regulations.

A defining characteristic of WP6 is its cross-sector orientation. As identified in the project's landscape analysis, diverse industry sectors often maintain varied cybersecurity practices, which can create inconsistencies and vulnerabilities. To address this, WP6 integrates requirements from three distinct critical sectors: Healthcare, Public Administration (Governance), and Academia.

Through Task 6.6 (Sector-Specific Cybersecurity Harmonization), the project seeks to bridge differences between these domains and advocate for the adoption of a consistent cybersecurity framework. By engaging with multiple industry sectors, this initiative ensures that there is a uniform set of cybersecurity standards and practices. This harmonization reduces vulnerabilities resulting from inconsistent practices and ensures that every sector is fortified against cyber threats.

Consequently, the compliance monitoring activities described in this deliverable are designed to be adaptable. They must support specific sectoral requirements (such as patient data privacy in healthcare or open network security in academia) while enforcing a harmonized baseline of security controls mandated by cross-cutting regulations like NIS2 and GDPR.

3.2 Real-Time Compliance Oversight

Task 6.7, titled "Real-time Compliance Oversight," is led by BYTE with the participation of MAG, IANUS, DSA, and KETYAK. The primary role of this task is to maintain the integrity of established cybersecurity standards by introducing tools that monitor and audit adherence in real time.

Unlike traditional compliance assessments which may be static or periodic, this task focuses on providing instant insights into how organizations are complying with established norms. It is designed to flag any deviations or non-compliance instantly. By providing real-time oversight, Task 6.7 ensures that standards are not just set but are also followed diligently. The capability for prompt detection ensures that deviations are rectified immediately, thereby minimizing potential security risks.

© CYberSynchorny, 2025

This task effectively operationalizes the governance frameworks defined in other tasks by deploying the technical means to verify them. It serves as the foundation for the compliance monitoring and auditing tools within the CYBERGOPLUS module, assessing entities' adherence to cybersecurity standards.

3.3 Dependencies with Tasks 6.1–6.6

Task 6.7, "Real-time Compliance Oversight," does not operate in isolation. It relies heavily on the outputs, frameworks, and data generated by the preceding and concurrent tasks within Work Package 6. While Tasks 6.1 through 6.6 focus on the creation of content, tools, and standards, Task 6.7 focuses on the verification and measurement of their adoption and effectiveness.

The specific dependencies are outlined below:

- **Dependence on Task 6.1 (Cyber Awareness Platform) and Task 6.4 (Immersive Gamified Cybersecurity):** These tasks aim to engage societal segments and utilize gamification for impactful learning. They provide the raw data required to measure "human-centric" compliance. Task 6.7 requires input from the Cyber Awareness Platform (T6.1) and Gamified Training Modules (T6.4) to verify metrics such as training completion rates and user engagement levels, ensuring that the objective of cultivating a secure culture is met.
- **Dependence on Task 6.2 (Proactive Social Engineering Assessments):** Task 6.2 involves targeted campaigns to simulate real-world threats like phishing to evaluate user responses. Task 6.7 relies on the data from these assessments to report on the organization's vulnerability to social engineering and to track improvements in staff resilience over time.
- **Dependence on Task 6.5 (Tailored Cybersecurity Policy Tools):** This task focuses on developing adaptable tools and templates to assist in the creation of cybersecurity policies. These policies form the "ruleset" that Task 6.7 monitors. Without the defined policies and regulations from T6.5 (e.g., Access Control Policy, Data Protection Policy), the compliance monitoring tools in T6.7 would lack the specific criteria against which to audit the organization.
- **Dependence on Task 6.6 (Sector-Specific Cybersecurity Harmonization):** Task 6.6 aims to bridge differences between diverse industry sectors (Healthcare, Public Administration, Academia) to advocate for a consistent cybersecurity framework. Task 6.7 depends on this harmonized framework to apply the correct audit parameters to specific entities. It ensures that the monitoring tools respect the unified standards established in T6.6 while verifying adherence across different sectors.

In summary, Tasks 6.1–6.6 establish the "what" (policies, training content, standards) and the "how" (training delivery, harmonization), while Task 6.7 establishes the "check" (monitoring, auditing, oversight) to ensure these elements are effectively implemented.

As the second iteration of the WP6 compliance-monitoring output, D6.8 aligns with the iteration versions of the WP6 deliverables it depends on, rather than their first versions. The Cybersecurity Policy Toolkit, whose 15 policy tools provide the governance evidence the framework monitors, is taken from its final version, D6.7 (Cybersecurity Policy Toolkit v2). D6.7 retains the 15 tools and their structure, refines them with organizational review templates, and adds the operationalization of those tools together with the CYberSynchrony core modules across the health, public-administration, and academia sectors, along with a roadmap toward the policy-compliance and certification indicators (KPI 9.1 and KPI 9.3). This creates a direct, two-way linkage: D6.7 operationalizes the policy tools per sector, and D6.8 monitors compliance against them, with the two deliverables addressing complementary indicators. The awareness and training evidence is aligned with the corresponding iteration deliverables, D6.5 (Cyber Awareness Platform v2) and D6.6 (Gamified Training Modules v2). D6.5 provides the pervasive awareness platform through which the policy toolkit and awareness services are delivered to users. D6.6 provides the gamified training modules, including the immersive cybersecurity game and the phishing-campaign exercises, extended in this iteration with integrated training on the EU legislative and policy framework. As D6.5 and D6.6 are being finalized in parallel with this deliverable, D6.8 aligns with their

© CYberSynchorny, 2025

defined scope and draws its awareness and training evidence from the capabilities they deliver, which are consistent across the iteration.

3.4 Deliverable Evolution and Key Outputs

Deliverable 6.8 is the second and final version of the Compliance Monitoring report for Task 6.7. It takes the framework established in D6.4 as its direct input and advances it from the definition and architecture stage to a fully specified operational design. Where v1 deliberately concentrated on defining the tools, architecture, communications and requirements necessary for real-time oversight, v2 builds on that foundation by specifying in concrete terms how the CYBERGOPLUS module consumes the outputs of the integrated components to produce compliance evidence, and by extending the framework with additional sources, a structured measurement methodology, and end-to-end operational detail. This section summarizes the foundations carried forward from v1 and the advances introduced in v2.

3.4.1 Foundations Carried Forward from v1 (Definition and Architecture)

The following outputs were established in v1 and are retained in full as the foundation of v2. They define the scope within which the present deliverable operates and are extended rather than revised.

- **Methodological Framework Definition**: The methodology for the CYBERGOPLUS module, detailing how compliance is monitored, audited and visualized, established in v1 and retained as the operating basis of v2.
- **Translation of Policy into Technical Requirements**: The mapping of the 15 static Policy Tools developed in D6.3 (Cybersecurity Policy Toolkit) to measurable technical requirements, including the mapping of policy clauses to specific digital indicators.
- **Technical Architecture Blueprint**: The technical architecture of the compliance module, specifying its integration with the Cyber Awareness Platform (D6.5) and the use of components such as Wazuh and Keycloak.
- **Initial Baseline Assessment**: The As-Is assessment of the pilot use cases, drawing on analysis from D2.2 and supplementary surveys, identifying current maturity gaps in compliance enforcement.

3.4.2 Advances Introduced in v2 (Specification, Integration and Validation Framework)

Building directly on the v1 foundations above, v2 advances the framework from a defined architecture to a fully specified design grounded in the concrete outputs of the integrated components. The advances delivered in this version are:

- **Detailed Specification of the Monitoring Engine**: The CYBERGOPLUS module is specified at the level of the actual evidence artifacts produced by the integrated components, defining precisely which output of each component is consumed to satisfy each control retained from v1.
- **Extension of the Evidence Base**: The evidence base established in v1 is extended with additional sources mapped to existing controls, including the threat-intelligence and secure-sharing outputs of MISP, the team trustworthiness and readiness assessment produced by the trustSense tool, and complementary enrichment from external threat-intelligence and analytics tooling. No new compliance controls are introduced.
- **KPI Measurement Methodology**: A structured methodology for compliance measurement is defined, specifying for each indicator its baseline, target population, measurement window, evidence source and associated confidence, and distinguishing preparatory indicators from outcome indicators.
- **End-to-End Compliance Workflows**: The component-level descriptions of v1 are advanced into concrete cross-module workflows that trace evidence from its generation in a component, through validation in CYBERGOPLUS, to the resulting compliance finding and remediation action.

© CYberSynchorny, 2025

- **Consolidated Per-Pilot Narratives**: The sector-specific requirements defined in v1 are consolidated into integrated, end-to-end narratives for each pilot (Healthcare, Public Administration, Academia), each presented in a single place rather than distributed across the document.
- **Ethics, Privacy and Data Protection**: The data collection underpinning compliance monitoring is assessed for privacy and data-protection impact, with equivalent scrutiny applied across all categories of evidence gathered.
- **Validation Framework**: A validation framework is defined for the operational assessment of the monitoring system within the pilot environments. In line with the project timeline, this version establishes the validation framework and the measurement methodology, while longitudinal operational validation within the pilots continues beyond this submission.

© CYberSynchorny, 2025

4 Compliance Framework Foundations

This chapter establishes the regulatory and theoretical foundations upon which the CYberSynchrony compliance monitoring mechanism is built. To ensure that the "Real-Time Oversight" capability defined in Task 6.7 is effective, it must be grounded in specific legal obligations, industry standards, and organizational policies rather than abstract metrics. This section outlines the regulatory baselines, primarily the NIS2 Directive and GDPR, that dictate the mandatory security controls for European essential entities.

Furthermore, this framework integrates the outputs of D6.3 (Cybersecurity Policy Toolkit), ensuring that the automated monitoring tools are configured to verify specific policy clauses defined in the project's standardized templates. The following sections detail how these regulatory requirements are harmonized with sector-specific needs for Healthcare, Public Administration, and Academia, creating a unified basis for the technical indicators described later in this document.

4.1.1 NIS2 Directive and National Implementation Frameworks

The NIS2 Directive (Directive (EU) 2022/2555) establishes cybersecurity risk management measures and reporting obligations for essential entities. To operationalize real-time compliance monitoring within the CYberSynchrony ecosystem, the project utilizes the security measures framework established by the Digital Security Authority (DSA) of Cyprus in Decision 389/2020. This framework structures the high-level requirements of NIS2 into measurable domains, including Governance, Protection, Defence, and Resilience. It is further aligned with the overall European approach for NIS2 implementation as defined by ENISA and the Member State national authorities[6].

The compliance monitoring approach translates these regulatory obligations into verifiable digital evidence. The CYBERGOPLUS module ingests specific data points from the platform's active components to verify that specific controls are operational, rather than relying on manual self-assessments.

Regulatory Control Mapping Operational Compliance Flow

To ensure traceability, the framework establishes a clear lineage from platform functionality to collected evidence, and finally to the satisfied regulatory requirement. The following table reconciles the CYberSynchrony platform capabilities with the specific evidence artifacts generated and the corresponding NIS2-aligned controls defined in the national framework and international standards.

Table 3: Mapping of Platform Functionality to Regulatory Evidence

CYberSynchrony Functionality (Tool)	Evidence Collected (Digital Artifact)	Source Deliverable	Regulatory Requirement Satisfied (DSA 389/2020 / ISO 27001 / NIST) [3][4][5]
Automated Asset Discovery (Wazuh)	Real-time inventory logs listing active IP addresses, hostnames, and OS versions.	D4.1	Asset Management (AM2) Inventory of assets (ISO A.8.1.1 / NIST CM-8)
Identity Management (Keycloak)	Configuration logs verifying MFA enforcement and Single Sign-On (SSO) usage.	D6.5	Identity & Access (IAM4) Strong authentication (ISO A.9.1.2 / NIST AC-3)

© CYberSynchorny, 2025

Network Traffic Analysis (Suricata)	EVE JSON logs confirming use of secure protocols (TLS 1.2/1.3) and absence of cleartext.	D4.1	Network Security (NS4) Secure communication protocols (ISO A.13 / NIST SC-8)
Vulnerability Detection (Wazuh)	CVE Reports detailing open Critical/High severity vulnerabilities on endpoints.	D4.1	Vulnerability Management (VM1) Vulnerability scanning (ISO A.12.6 / NIST RA-5)
Software Security Scanning (SNYK)	SAST reports identifying code-level vulnerabilities and dependency risks.	D6.1	Vulnerability Management (VM1) Vulnerability scanning (ISO A.12.6 / NIST RA-5)
Gamified Training (Cybergate)	User participation records and simulation outcome scores (e.g., Phishing defense).	D6.2	Awareness & Training (TA2) Information Security training (ISO A.7.2.2 / NIST AT-1)
Immutable Logging (Hyperledger)	Blockchain hash records proving logs and backups have not been altered.	D5.2, D4.1	Data Integrity (DS) Information integrity (NIS2 Art 21 / ISO A.8)
Policy Management (CYBERGOPLUS)	Digital policy versions and user acknowledgment timestamps.	D6.3	Governance (GOV3) Policies and procedures (ISO A.5.1 / NIST PL-1)
Threat Intelligence Enrichment (MISP)	Curated IoC events, attribute correlations, and STIX/IDS-rule exports supplied to the detection components and matched against telemetry.	D5.3	Incident & Event Management (EIM1) - Incident detection and handling (ISO A.16.1 / NIST IR-4)
Secure Threat-Information Sharing (MISP)	PGP/S-MIME signed and encrypted intelligence exchanges, sharing-group access records, and TLP taxonomy tags.	D5.3	Data Integrity (DS) - Secure information transfer (ISO A.13.2 / NIST AC-4)
Team Trustworthiness Assessment (trustSense)	Team trustworthiness and socio-technical risk values with recommended mitigation actions, derived from a structured team self-assessment.	D4.3[18]	Risk assessment methodology (ISO 27001 Cl. 6.1.2 / NIST RA-3)
External Threat-Intelligence Feeds (Kitsune)	Complementary dark-web and leak-monitoring feeds (malicious IP/URL/hash, credential-leak alerts) enriching the detection layer.	D3.3[19]	Incident & Event Management (EIM1) - Incident detection, complementary enrichment (ISO A.16.1 / NIST IR-4)

© CYberSynchorny, 2025

Evidence Requirements for Compliance Monitoring

The following subsections detail the evidence required to validate compliance for the core NIS2 domains. The monitoring engine relies on specific data streams to validate the controls listed above.

1. Asset and Configuration Evidence To satisfy Asset Management (AM) requirements, the system ingests dynamic data regarding the infrastructure state.

- **Functionality:** The Wazuh agent collects system-level telemetry including process executions and installed applications.
- **Evidence:** Automated inventory scans provided by D4.1 are reconciled against the authorized asset register defined in D2.2 requirements. This verifies control AM2 (Inventory of Assets).

2. Identity and Access Evidence Monitoring compliance with Identity and Access Management (IAM) focuses on authentication mechanisms.

- **Functionality:** Keycloak manages user identities and enforces Multi-Factor Authentication (MFA).
- **Evidence:** Authentication logs (success/failure) and configuration exports from D6.1 provide proof that strong authentication (IAM4) is active and that privileged access is restricted.

3. Technical Security Evidence To validate Network (NS) and System Security (SS), the platform ingests traffic and configuration data.

- **Functionality:** Suricata performs deep packet inspection to detect insecure protocols.
- **Evidence:** EVE JSON logs from D4.1 serve as evidence for Secure Communication Protocols (NS4), confirming that data in transit is encrypted.

4. Resilience and Awareness Evidence Verifying Resilience (BCR) and Awareness (TA) requires evidence of testing and human engagement.

- **Functionality:** The Cybergate platform simulates attack scenarios (e.g., Ransomware) to test staff response.
- **Evidence:** Training completion records and simulation scores from D6.2 provide objective evidence for Information Security Awareness (TA2) and verify that staff are prepared for incident response.

4.2 Sector-Specific Requirements (Healthcare, Governance, Academia)

While the CYberSynchrony platform applies a harmonized compliance framework, distinct sectors possess unique operational risks, asset types, and regulatory nuances that require tailored monitoring configurations. This section details the specific requirements for the Healthcare sector, as identified during the end-user analysis and pilot definition phases.

4.2.1 Healthcare Sector Requirements

The healthcare sector represents a critical domain within the NIS2 Directive scope, characterized by the processing of sensitive personal data and the reliance on life-critical systems. The digitization of health services, including Electronic Health Records (EHR) and Remote Patient Monitoring (RPM), has expanded the attack surface, making data privacy and system availability paramount.

Regulatory and Operational Context

Healthcare entities are classified as essential entities under NIS2 and must adhere to strict data protection standards defined by the GDPR, particularly regarding special category data (health data). The primary security objectives identified for this sector are Confidentiality (preventing unauthorized access to patient records), Integrity (ensuring medical data accuracy), and Availability (ensuring systems function during emergencies).

© CYberSynchorny, 2025

Specific Compliance Monitoring Requirements

To address the specific threats facing healthcare, such as ransomware and data breaches involved in IoMT (Internet of Medical Things), the compliance monitoring module must verify the following sector-specific controls:

- **Patient Data Confidentiality and Encryption:** Given the sensitivity of Protected Health Information (PHI), the monitoring system must verify that strong encryption is applied to patient databases and data in transit.
    - *Required Evidence:* Logs from the Data Integrity and Encryption module confirming that patient records stored or transferred via the Secure Data Exchange Platform utilise validated encryption protocols.
- **IoMT and Medical Device Security:** The proliferation of connected medical devices (pacemakers, insulin pumps) introduces vulnerabilities that traditional IT monitoring may miss.
    - *Required Evidence:* Network traffic analysis logs (Suricata) identifying communication patterns of IoMT devices to detect anomalies indicative of compromise, such as unauthorized external connections.
- **Availability and Ransomware Resilience:** Healthcare operations require continuous up time. Ransomware poses a direct threat to patient safety by locking critical systems.
    - *Required Evidence:* Backup verification logs (from the Asset Management Policy requirements) confirming that critical medical data is backed up regularly and that restoration tests are successful.
- **Access Control for Medical Personnel:** Access to health records must be restricted based on clinical roles (e.g., doctors, nurses, administrators).
    - *Required Evidence:* Identity Management logs (Keycloak) verifying that Multi-Factor Authentication (MFA) is enforced for all staff accessing EHR systems and that Role-Based Access Control (RBAC) policies are active.

Table 4: Healthcare-Specific Compliance Indicators

Specific Risk / Asset	Monitoring Requirement	Evidence Source	Regulatory Alignment
Electronic Health Records (EHR)	Verify encryption at rest and strict RBAC enforcement.	D4.1 (Wazuh/Sysmon), D6.1 (Keycloak)	GDPR Art. 32, NIS2 Art. 21
Internet of Medical Things (IoMT)	Monitor for unauthorized firmware changes or traffic anomalies.	D4.1 (Suricata/Anomaly Detector)	NIS2 Supply Chain Security
Telehealth Services	Ensure secure communication channels (TLS 1.3) for remote consultations.	D4.1 (Suricata EVE JSON)	HIPAA (Reference), GDPR
Medical Staff Access	Audit logs for unauthorized access attempts to high-profile patient data.	D6.1 (Keycloak), D4.1 (Wazuh)	ISO 27001 A.9, GDPR

© CYberSynchorny, 2025

These specific requirements inform the configuration of the CYBERGOPLUS module for the Healthcare pilot (Use Case 1), ensuring that the generic technical controls are applied with the necessary rigor for medical environments.

4.2.2 Public Administration Sector Requirements

The Public Administration sector, encompassing local government and municipalities, faces distinct challenges related to the provision of essential services, the management of public records, and the protection of citizen data. As identified in the CyberSynchrony use case analysis, this sector is characterized by a high reliance on legacy systems, a broad attack surface due to e-governance platforms, and strict accountability requirements.

Regulatory and Operational Context

Public administration entities are classified as essential under the NIS2 Directive, necessitating robust measures to ensure service continuity and data integrity. The primary security objectives for this sector include preventing the disruption of critical services (e.g., water supply, emergency response), safeguarding sensitive citizen information against data breaches, and maintaining trust in digital public infrastructure.

Specific Compliance Monitoring Requirements

To address threats such as ransomware, DDoS attacks, and supply chain compromises, the compliance monitoring module must verify the following sector-specific controls:

- **Service Availability and Resilience:** Given the critical nature of municipal services, the monitoring system must verify that e-governance platforms and utility control systems maintain high availability.
    - *Required Evidence:* Availability logs from the CYBRITE module (Wazuh/Suricata) detecting DDoS attempts and confirming the uptime of critical public-facing portals.
- **Integrity of Public Records:** The unauthorized alteration of public records (e.g., land registry, vital statistics) undermines governance. Monitoring must ensure data immutability.
    - *Required Evidence:* Validation logs from the CROSS-CORE Evidence Integrity System (Blockchain), providing cryptographic proof that stored public records have not been tampered with.
- **Secure Digital Identity Management:** With the shift to digital IDs for citizen services, the system must verify robust authentication mechanisms to prevent identity theft and unauthorized access.
    - *Required Evidence:* Identity Management logs (Keycloak) confirming the enforcement of strong authentication protocols for accessing digital service portals.
- **Supply Chain Security:** Public administrations often rely on numerous third-party vendors for software and infrastructure. Monitoring must track the security posture of these dependencies.
    - *Required Evidence:* Supply chain risk assessment reports (from CYBERRA/SNYK) identifying vulnerabilities in third-party software components integrated into municipal systems.

Table 5: Public Administration-Specific Compliance Indicators

Specific Risk / Asset	Monitoring Requirement	Evidence Source	Regulatory Alignment
E-Governance Portals	Monitor for DDoS signatures and service uptime.	D4.1 (Suricata), D4.1 (Wazuh)	NIS2 Art. 21 (Business Continuity)
Public Records Databases	Verify data integrity via blockchain hashing.	D5.2 (Evidence Integrity System)	GDPR (Integrity), NIS2

© CYberSynchorny, 2025

Digital Citizen IDs	Audit authentication logs for anomaly detection.	D6.1 (Keycloak), D4.1 (Anomaly Detector)	eIDAS, GDPR
Utility Control Systems (OT)	Monitor traffic at IT/OT boundaries for unauthorized access.	D4.1 (Suricata/Wazuh)	NIS2 (Supply Chain & Network Security)

These requirements guide the configuration of the CYBERGOPLUS module for the Public Administration pilot (Use Case 2), ensuring that compliance monitoring addresses the specific needs of local governance and digital public services. Specifically, in alignment with the European digital identity framework, the monitoring module validates compliance with the **eIDAS Regulation**, specifically regarding trust services and electronic identification. For Public Administration (UC2) and Academia (UC3), the system monitors the integrity of digital signatures and authentication timestamps. This ensures that the "Continuous Threat Detection" principle extends to the validity of digital transactions, a critical requirement for cross-border services within the Digital Single Market.

4.2.3 Academia Sector Requirements

The Academia sector operates under a unique security paradigm characterized by "open by default" networks, decentralized IT management, and a high volume of Bring Your Own Device (BYOD) usage. As defined in Use Case 3 (Cybersecure University), this sector faces significant threats regarding Intellectual Property (IP) theft, ransomware attacks on academic repositories, and the compromise of student personal data.

Regulatory and Operational Context

While Academic institutions must comply with GDPR regarding student and staff data, they also face specific requirements related to the protection of academic integrity and funding grant stipulations. The security objective here is to balance the academic freedom and open information exchange required for academic work with the stringent controls needed to protect high-value IP and sensitive personal records.

Specific Compliance Monitoring Requirements

To address the specific vulnerabilities of open campus networks and decentralized academic labs, the compliance monitoring module must verify the following sector-specific controls:

- **Academic Data Integrity and Protection:** Research institutions must ensure that critical research data is not altered or stolen. Monitoring must verify that access to sensitive research segments is restricted and that data integrity is maintained.
    - *Required Evidence:* Integrity verification logs from the Evidence Integrity System (Blockchain) confirming that critical research datasets remain unaltered, alongside access logs (Wazuh) showing restricted entry to IP repositories.
- **Student Record Privacy:** Universities handle vast amounts of Personally Identifiable Information (PII). Monitoring must ensure these records are encrypted and accessed only by authorized administrative staff.
    - *Required Evidence:* Database encryption status logs and Identity Management audits (Keycloak) verifying that Role-Based Access Control (RBAC) separates student records from general faculty access.
- **E-Learning Platform Availability:** The reliability of online learning platforms is critical for continuity. These platforms are frequent targets for DDoS attacks.
    - *Required Evidence:* Network traffic analysis (Suricata) and server health logs (Sysmon/Wazuh) monitoring the availability and latency of Learning Management Systems (LMS).

© CYberSynchorny, 2025

- **Open Network Security (BYOD):** Given the prevalence of personal devices, the system must monitor for rogue devices and anomalous traffic patterns originating from within the campus network.
    - *Required Evidence:* Network anomaly detection logs (Anomaly Detector IDS) identifying suspicious lateral movement or unauthorized scanning activities within open campus Wi-Fi segments.

Table 6: Academia-Specific Compliance Indicators

Specific Risk / Asset	Monitoring Requirement	Evidence Source	Regulatory Alignment
Academic Repositories	Verify data integrity hashing and access restrictions.	D5.2 (Evidence Integrity), D4.1 (Wazuh)	IP Protection Laws, Funding Agreements
Student Information Systems	Audit access logs for PII and verify encryption.	D4.1 (Wazuh), D6.1 (Keycloak)	GDPR (Confidentiality)
E-Learning/LMS	Monitor platform uptime and DDoS signatures.	D4.1 (Suricata), D4.1 (Sysmon)	NIS2 (Business Continuity)
Campus Wi-Fi (BYOD)	Detect anomalous traffic flows and rogue access points.	D4.1 (Anomaly Detector/Suricata)	Network Security Best Practices

These requirements define the configuration of the CYBERGOPLUS module for the University pilot (Use Case 3), ensuring that the monitoring framework accommodates the open nature of academic environments while securing critical assets. The operational validation of the CYBERGOPLUS module will be conducted through the cross-sector Use Cases defined in WP8. The tool will be deployed in three distinct pilot environments to demonstrate its adaptability:

- **Use Case 1 (Healthcare):** The module will be integrated with the Hospital Information System (HIS) to validate the "Asset-Centric Protection" principle, specifically monitoring IoMT devices and patient data access (Tool 10).
- **Use Case 2 (Public Administration):** The pilot will focus on the "Continuous Threat Detection" principle, validating the system's ability to detect threats against e-governance portals (Tool 8).
- **Use Case 3 (Academia):** The deployment will test the "Resilience and Recovery" principle, ensuring that open research networks can maintain continuity during simulated attacks (Tool 15).

4.3 Harmonised Cross-Sector Principles

While the specific operational contexts of Healthcare, Public Administration, and Academia differ, the fundamental requirements for cybersecurity compliance under NIS2 and ISO 27001 share a common logical foundation. To ensure the CYberSynchrony platform can scale across these diverse sectors without requiring bespoke development for every new use case, the compliance monitoring framework is built upon four Harmonised Cross-Sector Principles. These principles distil the regulatory obligations into functional categories that apply universally to all essential entities.

These principles serve as the bridge between the high-level legal text (e.g., NIS2 Article 21) and the low-level technical indicators collected by the platform components. Each principle is mapped to specific controls from the DSA 389/2020 framework to ensure regulatory alignment, and then to specific CYberSynchrony functionalities to ensure technical feasibility.

© CYberSynchorny, 2025

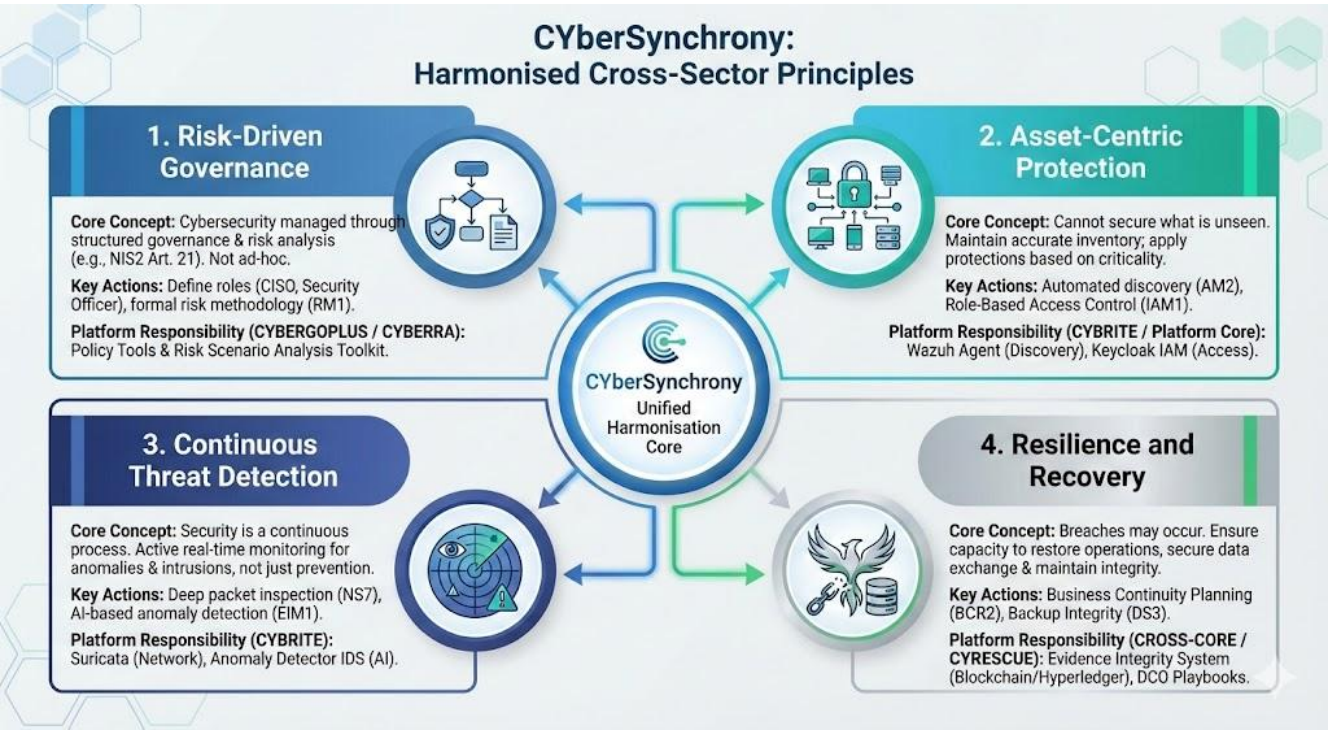

Figure 1: Harmonised Cross-Sector Compliance Principles mapping regulatory concepts to CYberSynchrony platform responsibilities.

4.3.1 Principle 1: Risk-Driven Governance

The first principle mandates that cybersecurity must be managed through a structured governance framework driven by risk analysis rather than ad-hoc decisions. This aligns with the NIS2 requirement for management bodies to approve cybersecurity measures and oversee their implementation.

- **Regulatory Control (DSA 389/2020):**
    - **GOV1 (Roles and Responsibilities):** Defines the obligation to assign specific security roles within the organization.
    - **RM1 (Risk Management Methodology):** Requires a formal methodology for identifying and treating risks.
- **CYberSynchrony Functionality:** The **CYBERGOPLUS** module operationalizes this principle. It provides the "Roles & Responsibilities Matrix" feature described in the requirements analysis, allowing organizations to digitally assign accountability. Furthermore, the **CYBERRA** module supports this by offering the "Risk Scenario Analysis Toolkit," which digitizes the risk assessment process.
- **Monitoring Evidence:** Compliance is verified by checking for the existence of active user accounts with "Admin" or "Security Officer" roles in the Identity Manager and the completion timestamps of digital risk assessment reports.

4.3.2 Principle 2: Asset-Centric Protection

This principle establishes that an organization cannot secure what it cannot see. It mandates the maintenance of an accurate inventory of all digital assets and the application of protective measures such as updates and access controls based on asset criticality.

- **Regulatory Control (DSA 389/2020):**

© CYberSynchorny, 2025

- **AM2 (Inventory of Assets):** Requires the identification and recording of all assets and their ownership.
- **IAM1 (Role-Based Access Control):** Mandates that access to these assets is restricted based on user roles.

- **CYberSynchrony Functionality:** The **CYBRITE** module, specifically through the **Wazuh** component, fulfills this principle by performing automated asset discovery and inventory management. Simultaneously, the **Keycloak Identity Manager** enforces access control policies, ensuring that only authorized users interact with the discovered assets.
- **Monitoring Evidence:** Compliance is monitored by reconciling the list of active endpoints detected by Wazuh agents against the authorized asset list and verifying authentication logs for proper access enforcement.

4.3.3 Principle 3: Continuous Threat Detection

Security is not a static state but a continuous process. This principle requires organizations to implement active monitoring capabilities to detect anomalies, intrusions, and potential incidents in real time, rather than relying solely on preventive measures.

- **Regulatory Control (DSA 389/2020):**
  - **NS7 (Intrusion Detection):** Requires measures to detect and prevent unauthorized access to the network.
  - **EIM1 (Event Readiness and Detection):** Mandates the capability to detect events that could threaten security.
- **CYberSynchrony Functionality:** The platform addresses this through the **Integrated Cybersecurity Monitoring Suite** (D4.1). **Suricata** performs deep packet inspection to identify network threats, whilst the **Anomaly Detector IDS** utilises AI-based logic to flag deviations in traffic patterns that may indicate novel attacks.
- **Monitoring Evidence:** Compliance is evidenced by the continuous stream of EVE JSON logs and alert records generated by these components, proving that the detection mechanisms are active and functioning.

4.3.4 Principle 4: Resilience and Recovery

The final principle acknowledges that breaches may occur and mandates that organizations must have the capacity to restore operations. This covers business continuity planning, backup management, and the integrity of data required for recovery.

- **Regulatory Control (DSA 389/2020):**
  - **BCR2 (Business Continuity Plan):** Requires a plan to maintain critical functions during a crisis.
  - **DS3 (Backup and Recovery):** Mandates regular backups and the ability to restore information.
- **CYberSynchrony Functionality:** The **CROSS-CORE** module supports resilience by facilitating secure data exchange and coordination during incidents . Additionally, the **Evidence Integrity System** (Hyperledger) ensures that the logs and backup hashes required for recovery have not been tampered with, ensuring a trusted restoration point.
- **Monitoring Evidence:** Compliance is verified through the "Evidence Integrity System" logs which validate the immutability of critical records, and through the tracking of incident response activities within the **Defensive Cyber Operations (DCO)** playbooks.

© CYberSynchorny, 2025

4.4 Mapping of Principles to CYberSynchrony Responsibilities

The following table summarizes the harmonization of these principles, mapping the regulatory requirement to the specific technical component within the CYberSynchrony ecosystem that assumes responsibility for fulfilling it.

Table 7: Mapping of Harmonised Principles to Platform Responsibilities

Harmonised Principle	Regulatory Reference (DSA 389/2020)	CYberSynchrony Module Responsibility	Key Component
1. Risk-Driven Governance	GOV1, RM1	CYBERGOPLUS / CYBERRA	Policy Tools, Risk Assessment Toolkit
2. Asset-Centric Protection	GOV1, RM1	CYBRITE / Platform Core	Wazuh Agent, Keycloak IAM
3. Continuous Threat Detection	NS7, EIM1	CYBRITE	Suricata, Anomaly Detector IDS
4. Resilience and Recovery	BCR2, DS3	CROSS-CORE / CYRESCUE	Evidence Integrity System, DCO Playbooks

Specifically, the principles are operationalized through the 15 policy tools defined in the Cybersecurity Policy Toolkit (D6.3). This alignment ensures that every technical principle monitored by the system is backed by a formal governance policy:

- **Principle 1 (Risk-Driven Governance):** Operationalized by **Tool 1** (Information Security Strategy), **Tool 2** (Information Security Governance Policy), and **Tool 3** (Risk Management Policy & Methodology).
- **Principle 2 (Asset-Centric Protection):** Operationalized by **Tool 4** (Asset Management Policy), **Tool 5** (Supplier and Third-Party Management Policy), **Tool 6** (Access Control Policy), **Tool 7** (Vulnerability and Patch Management Policy), **Tool 8** (Network Security Policy), **Tool 9** (Systems & Applications Security Policy), and **Tool 10** (Data Protection Policy).
- **Principle 3 (Continuous Threat Detection):** Operationalized by **Tool 3** (Risk Management Policy & Methodology), **Tool 12** (Human Resources Security, Training & Awareness Policy), **Tool 13** (Physical Security Policy), and **Tool 14** (Incident and Event Management Policy).
- **Principle 4 (Resilience and Recovery):** Operationalized by **Tool 15** (Business Continuity and Disaster Recovery Policy).

The mapping presented above serves as the operational "bridge" between the abstract regulatory requirements of the NIS2 Directive and the concrete technical configurations monitored by the CYBERGOPLUS module. While the Harmonised Principles defined in Section 4.3 provide the high-level logic for compliance, categorizing obligations into Governance, Protection, Detection, and Resilience, it is the **Cybersecurity Policy Toolkit (D6.3)** that translates these concepts into enforceable organizational rules.

This structure ensures a complete chain of custody for compliance evidence. By operationalizing **Principle 1 (Risk-Driven Governance)** through **Tools 1–3**, the framework ensures that security is not treated as an ad-hoc technical task but as a board-level strategic imperative, directly satisfying the "Management Responsibility"

© CYberSynchorny, 2025

requirements of NIS2 Article 20. Similarly, the heavy concentration of tools under **Principle 2 (Asset-Centric Protection** spanning **Tools 4 through 10** reflects the operational reality that the majority of auditable controls are preventive in nature, focusing on the correct configuration of assets, identities, and networks.

Furthermore, the explicit mapping of **Principle 3 (Continuous Threat Detection)** to **Tools 12–14** shifts the focus from purely preventive security to the "active defense" posture required by modern regulations. It mandates that organizations not only possess security policies but also active mechanisms for Incident Management and Assurance Auditing. Finally, **Principle 4 (Resilience and Recovery)**, anchored by **Tool 15**, closes the loop by ensuring that when preventive and detective measures fail, a formally defined Business Continuity plan is available to restore operations. The table below indicates the complete list of controls adopted by DSA and used for compliance monitoring:

Table 8: DSA Derived Controls and Description

Control	Title	Description
GOV1	Roles and Responsibilities	Defines the obligation to assign specific cybersecurity roles and accountability within the organization, such as a NIS Officer, DPO, or Risk Manager.
GOV3	Policies and Procedures	Requires formally adopted security policies and procedures to be defined, approved, and kept current, operationalized through the 15 policy tools of D6.3.
RM1	Risk Management Methodology	Requires a formal methodology for identifying, assessing, and treating security risks, reviewed at least annually (ISO 27001 Cl. 6.1.2 / NIST RA-3).
AM1	Asset Management	Requires assets to be managed throughout their lifecycle, with protective measures such as updates and access controls applied according to asset criticality.
AM2	Inventory of Assets	Requires the identification and recording of all digital assets and their ownership.
IAM1	Role-Based Access Control	Mandates that access to assets is restricted based on defined user roles.
IAM4	Strong Authentication	Requires strong, multi-factor authentication for access, particularly for privileged accounts.
NS4	Secure Communication Protocols	Requires network communications to use secure, encrypted protocols, for example TLS 1.2 or higher.
NS7	Intrusion Detection	Requires measures to detect and prevent unauthorized access to the network.
VM1	Vulnerability Scanning	Requires regular scanning for vulnerabilities and the risk-based remediation of identified weaknesses.
EIM1	Event Readiness and Detection	Mandates the capability to detect events and anomalies that could threaten security in real time.
EIM2	Incident Analysis and Response	Requires the capability to analyze, handle, and respond to detected incidents, with an auditable record, aligned with NIS2 Article 21 incident handling.

© CYberSynchorny, 2025

BCR2	Business Continuity Plan	Requires a plan to maintain critical functions during a crisis or disruption.
DS3	Backup and Recovery	Mandates regular backups and the ability to restore information.
DS1	Data Lifecycle Management	Requires that data is protected from unauthorized change, access, or tampering
TA2	Awareness and Training	Requires cybersecurity awareness and training for staff and users, with evidence of engagement.

© CYberSynchorny, 2025

5 CYberSynchrony Compliance Monitoring Approach (High-Level)

This chapter outlines the operational methodology used to implement the "Real-Time Compliance Oversight" capability defined in Task 6.7. Building upon the regulatory foundations and harmonized principles established in the previous chapters, this section defines how the CYberSynchrony platform technically achieves continuous auditing. The approach moves away from static, document-based compliance towards a dynamic model where technical evidence serves as the primary validation mechanism. This methodology is designed to be agnostic of the specific underlying technology of the end-user while remaining rigorous enough to satisfy the requirements of the NIS2 Directive and ISO 27001 standards. In this version, the approach is specified against the concrete evidence artifacts produced by the integrated components, the measurement methodology for the associated KPIs is defined explicitly, and the evidence base is extended with the additional sources introduced in Chapter 4.

5.1 Objectives of Compliance Monitoring

The primary objective of the compliance monitoring framework is to provide stakeholders with an accurate, real-time view of their cybersecurity posture, minimizing the latency between a non-compliant event and its remediation. Specifically, the monitoring approach aims to achieve the following four operational objectives:

- **Enable Real-Time Visibility:** To transition from periodic audits (annual or quarterly) to continuous monitoring, providing instant alerts when a specific control fails or a policy is violated.
- **Automate Evidence Collection:** To reduce the administrative burden on cybersecurity staff by automating the gathering of technical logs, configurations, and artifacts required for audits, addressing the resource constraints identified in the end-user analysis.
- **Ensure Cross-Sector Consistency:** To apply a unified auditing logic that allows for comparable compliance scoring across different sectors (Healthcare, Public Administration, Academia), facilitating the "Harmonization" objective of Work Package 6.
- **Support Risk-Based Prioritization:** To link compliance findings directly to risk scenarios, ensuring that organizations prioritize the remediation of non-compliance issues that pose the greatest threat to their critical assets.

5.2 Compliance Domains and Assessment Areas

The core purpose of defining specific compliance domains within the CYberSynchrony platform is to translate granular technical telemetry into high-level reporting that is understandable to management boards and non-technical decision-makers. As identified in the end-user requirement analysis (D2.2), a critical gap exists in generating "Management/Board Reporting" that provides a clear view of compliance status without requiring manual aggregation of technical logs.

It is important to note that the CYberSynchrony tool produces compliance data specifically derived from the active components integrated into the ecosystem (e.g., Wazuh, Suricata, Keycloak, SNYK). The platform does not promise a "full" audit of every administrative or physical control within an organization; rather, it provides continuous validation for the specific digital controls that can be monitored via the project's technical architecture.

The assessment areas are structured to generate reports like the survey outcomes analyzed in D2.2, specifically addressing the need for standardized content elements such as risk status, initiative progress, and compliance gaps. By mapping the technical evidence described in Section 4.1.1 to the policy domains defined in D6.3, the platform assesses the following specific areas:

- **Governance & Strategy:** Assessments focus on the assignment of security roles and the verification of policy acknowledgments by staff, aligned with the *Information Security Governance Policy* (Tool 2).

© CYberSynchorny, 2025

- **Asset Security:** Monitoring focuses on the reconciliation of discovered assets against authorized inventories and the verification of maintenance schedules, aligned with the *Asset Management Policy* (Tool 4).
- **Identity Control:** Assessments track the enforcement of Multi-Factor Authentication (MFA) and the management of privileged accounts, aligned with the *Access Control Policy* (Tool 6).
- **Threat Protection:** Monitoring evaluates the active status of anti-malware agents, firewall configurations, and intrusion detection systems, aligned with the *Systems & Applications Security Policy* (Tool 9) and *Network Security Policy* (Tool 8).
- **Vulnerability Status:** Assessments report on the presence of unpatched critical vulnerabilities and software dependencies, aligned with the *Vulnerability and Patch Management Policy* (Tool 7).
- **Resilience:** Monitoring verifies the successful execution of backup routines and the completion of business continuity drills, aligned with the *Business Continuity and Disaster Recovery Policy* (Tool 15).

5.3 Indicators, Metrics and Maturity Levels

To transition compliance from a qualitative assessment to a quantitative science, the CYberSynchrony platform utilizes a hierarchical measurement model. This model aggregates low-level technical signals (Indicators) into operational performance measures (Metrics) and finally into high-level strategic targets (KPIs). This structure ensures that the technical evidence collected by components like Wazuh and Suricata directly supports the validation of the project's Grant Agreement objectives.

5.3.1 Key Compliance Indicators (KCIs)

Key Compliance Indicators (KCIs) serve as the atomic units of the monitoring framework. A KCI represents a specific, binary check performed against a technical asset to verify a single control requirement from the DSA 389/2020 framework. These indicators are derived directly from the raw telemetry produced by the platform's security tools.

The system utilizes the following logic for KCI generation:

- **Input:** Raw Log (e.g., Suricata EVE JSON, Wazuh Syslog).
- **Logic:** Boolean evaluation against a policy threshold (e.g., tls_version >= 1.2).
- **Output:** Pass/Fail flag with timestamp and Asset ID.

In this version, the generation logic is specified across the full set of integrated components and their documented outputs. Every check follows the same atomic form and yields the same Pass/Fail result with a timestamp and Asset ID. The following table enumerates a representative KCI for each evidence source, expressed as the boolean check it performs and the control it verifies. The list is representative rather than exhaustive, and the complementary sources (MISP, Kitsune, trustSense) enrich existing controls without introducing new ones.

Table 9: Operational Compliance Indicators by Evidence Source

Evidence Source	Source Output Artifact	Representative Check (KCI Logic)	Control Verified
Wazuh	syscollector inventory logs (JSON)	Active asset is registered in the authorized inventory	AM2
Wazuh	Vulnerability Detector CVE reports	No open Critical or High CVE remains on the asset	VM1

© CYberSynchorny, 2025

Wazuh	Security Configuration Assessment (SCA) results	Hardening-check pass rate meets or exceeds the threshold	GOV3
Wazuh	File Integrity Monitoring (FIM) events	FIM is enabled and reporting on critical system paths	DS
Suricata	EVE JSON (TLS records)	tls.version is 1.2 or higher on monitored flows	NS4
Suricata	EVE JSON (alert records)	The sensor is producing alerts within the monitoring window	NS7
Keycloak	Realm settings (JSON)	MFA is set to Required for privileged user groups	IAM4
Keycloak	Authentication event logs (type=LOGIN, mfa)	Privileged logins are performed with MFA enabled	IAM4
SNYK	SAST and SCA findings	No unresolved Critical or High finding remains on the project	VM1
AI-based IDS / Anomaly Detector	Anomaly detection records with score	No anomaly above the score threshold remains unhandled	EIM1
MISP	IoC events and feeds (JSON) with TLP tags	The threat-intelligence feed is current and shared data carries a TLP tag	EIM1
MISP	Signed and encrypted exchange records	Intelligence exports carry a valid PGP or S/MIME signature	DS
Kitsune	Credential-leak alerts	No unresolved credential-leak alert exists for the organization's domains	EIM1
Kitsune	Structured incidents with severity and IR status	High-severity incidents are closed within the SLA window	EIM2
DCO Playbooks	Playbook execution logs	A playbook executed and completed for each triggered incident	EIM2
Cybergate	Mission completion logs and scores	Mandatory training modules are completed at or above the score threshold	TA2
trustSense	Team trustworthiness score (out of 5)	The team assessment is current and the score meets the threshold	RM1
Evidence Integrity System	Blockchain transaction hash	The archive hash is valid and unaltered	DS3

5.3.2 Mapping Controls to KPIs

While the atomic KCIs enumerated in the preceding table establish whether a single control passes or fails on a given asset, this section aggregates those indicators into the operational metrics and strategic KPIs that the framework reports. The measurement hierarchy is therefore three-tiered: individual KCIs (Pass/Fail) are

© CYberSynchorny, 2025

aggregated into Metrics (rates, ratios and latencies), which in turn evidence the high-level KPIs defined in the CYberSynchrony Description of Action (DoA). The following table presents this mapping of regulatory controls to specific digital evidence, the associated measurement logic, and the linked DoA KPIs. It reconciles the operational requirements identified in D2.2 (KPIRs) with the monitoring capabilities established in D4.1, and in this version it is extended to cover additional controls and to incorporate the evidence sources introduced in Chapter 4.

Table 10: Comprehensive Compliance Indicator Map

Regulatory Control (DSA 389/2020)	CYberSynchrony Evidence Source	Technical Evidence Artifact	Metric / Measurement Logic
AM2: Inventory of Assets	Wazuh (D4.1)	syscollector inventory logs (JSON)	Unmanaged Asset Ratio: Percentage of active IPs detected by Suricata but not registered in Wazuh inventory.
VM1: Vulnerability Scanning	Wazuh / SNYK (D4.1/D6.1)	CVE Vulnerability Reports (Severity: Critical/High)	Mean Time to Remediate (MTTR): Average time between CVE detection timestamp and patch application timestamp.
IAM4: Strong Authentication	Keycloak (D6.1)	Authentication Event Logs (type=LOGIN, mfa=true/false)	MFA Coverage: Percentage of successful logins performed with Multi-Factor Authentication enabled.
EIM1: Incident Detection	Suricata / Anomaly Detector (D4.1), enriched by MISP (D5.3) and Kitsune	Alert Logs (event_type: alert, severity: 1); IoC feeds; credential-leak alerts	Detection Latency: Time delta between packet ingress timestamp and alert generation timestamp.
EIM2: Incident Analysis	DCO Playbooks (D4.1), Kitsune	Incident Ticket Status Changes (Open -> Closed)	Resolution Rate: Percentage of high-severity incidents closed within the SLA window (e.g., 24 hours).
TA2: Awareness Training	Cybergate (D6.2)	User Mission Completion Logs & Scores	Training Coverage: Percentage of active directory users who have completed mandatory simulation modules.
DS3: Backup and Recovery	Evidence Integrity System (D5.2 / D5.4)	Blockchain Transaction Hash (Hyperledger)	Integrity Verification Rate: Percentage of backup archives with valid, unaltered blockchain hashes.
NS4: Secure Communication Protocols	Suricata (D4.1)	EVE JSON (TLS records)	Secure Protocol Coverage: Percentage of monitored flows using TLS 1.2 or higher.

© CYberSynchorny, 2025

NS7: Intrusion Detection	Suricata / Anomaly Detector (D4.1)	Alert logs and anomaly detection records	Detection Coverage: Percentage of monitored network segments with an active, alerting sensor.
GOV3: Policies and Procedures	CYBERGOPLUS (D6.3)	Digital policy versions and adoption records	Policy Adoption Rate: Percentage of the 15 D6.3 policy tools with current, formally adopted digital records.
RM1: Risk Management Methodology	trustSense; Risk Scenario Analysis Toolkit (D3.2)	Team trustworthiness scores (out of 5); Risk Register	Risk Assessment Currency: Percentage of teams and critical assets with a current risk assessment.

5.3.3 Maturity Levels

To translate these metrics into a management-friendly format, the platform aggregates KCIs into a five-level maturity scale. This scale aligns with the "Cybersecurity Maturity Controls" decision (K.D.P. 245/2024) referenced in the Risk Management Policy. In this version the derivation is made explicit: the maturity level for a given assessment area is obtained from the aggregate pass rate of the KCIs that populate it (the proportion of passing checks over the total evaluated) and then banded into the five levels below. An overall organizational maturity is obtained by aggregating across the six assessment areas defined in Section 5.2. So that a high average cannot mask a critical weakness, failures of higher-criticality controls are surfaced separately rather than diluted into the aggregate.

- **Initial (0-20% Compliance)**: Controls are non-existent or ad-hoc. Evidence shows high rates of unmanaged assets and open critical vulnerabilities, with little or no automated checking in place.
- **Managed (21-40% Compliance)**: Controls are documented but implementation is partial. MFA is active for administrators but not general staff, and Security Configuration Assessment runs but with a high proportion of failing checks.
- **Defined (41-60% Compliance)**: Controls are standardized. Automated scanning (VM1) is active and File Integrity Monitoring is enabled on critical paths, but remediation times (MTTR) exceed targets.
- **Quantitatively Managed (61-80% Compliance)**: Controls are measured and controlled. KPIs such as Incident Detection Time (KPI 2.2) are consistently met, and threat-intelligence and credential-leak enrichment are actively monitored.
- **Optimizing (81-100% Compliance)**: Focus on continuous improvement. Advanced anomaly detection (EIM1) is active and integrated with automated response playbooks, and high-severity incidents are consistently resolved within the SLA window.

Because each level is anchored to measurable KCI evidence rather than to self-assessment, an organization's position on the scale is reproducible, and its progression is actionable: advancing a level requires raising the pass rate of the specific failing controls identified within each assessment area. This provides the measurable progression path that the v1 baseline indicated was missing, where high policy adoption coexisted with low technical enforcement. The maturity level is reported both per assessment area and as an aggregate posture, and serves as the primary input to the board-level reporting defined in Chapter 6.

5.3.4 KPI Measurement Methodology

To ensure that every reported KPI is interpretable and verifiable, this version specifies a common measurement methodology. Each KPI is documented with five attributes: a baseline (the as-is value or condition from which change is measured), a target population (the entities or assets to which it applies), a measurement window (the

© CYberSynchorny, 2025

period and cadence of measurement), an evidence source (the artifact or dataset from which the value is computed), and a statement of confidence and limitations. KPIs are further classified as preparatory indicators, which evidence that the capability to measure an outcome has been put in place, or outcome indicators, which report a measured result. The operational KPIs in Table 9 follow this template, with their values computed from the KCIs described above. This section specifies in full the outcome KPI for which D6.8 is the designated methodological home, KPI 9.2, while the measured result is reported in D8.3.

KPI 9.2: Industry-specific Standard Adoption

- **Definition**: The adoption rate of industry-specific cybersecurity standards and best practices by NIS2-covered critical sectors. The KPI is owned by IANUS under Objective 9 and Task 6.6.
- **Type**: Outcome indicator, distinct from the operational indicators derived from runtime telemetry.
- **Target Population**: Entities in the NIS2-covered critical sectors, operationalized within CYberSynchrony through the three use-case participants representing Healthcare (UC1), Public Administration (UC2) and Academia (UC3).
- **Baseline**: No single official EU dataset reports a clean, sector-by-sector standard-adoption rate, so the baseline is triangulated from recognized proxy indicators. The closest direct proxy is the maintenance of formal, documented ICT-security practices, which Eurostat (2024) records for 57.87% of large enterprises, the size class of the use-case entities; ENISA's cross-sector maturity work is consistent with this, placing healthcare at upper-moderate maturity and public administration among the least mature. **Consolidating these indicators, the baseline adoption rate is taken as approximately 58%.** No core NIS2 sector figure is available for the academia use case, so its contribution is established from the project's own assessment. The consolidated baseline evidence and sources are presented in Annex II.
- **Measurement Window**: The baseline is established at the start of the measurement period through the project assessment; the post-implementation value is measured subsequently and reported in D8.3.
- **Evidence Source**: A survey and assessment of the adoption of industry-specific standards and best practices by the use-case participants, complemented by the external proxy indicators above for sectoral context.
- **Target**: An increase in the adoption rate to 65%. This target is operationalized through use-case adoption of the CYberSynchrony platform: adoption by two of the three use-case participants corresponds to 66.6%, which meets and justifies the 65% target.
- **Confidence and Limitations**: Three limitations apply. First, EU-level evidence is uneven and largely proxy-based rather than a direct compliance count. Second, the operational target population is small (three use cases), so the 66.6% figure is necessarily coarse-grained. Third, the external baseline figures are drawn from secondary sources and should be read as contextual proxies rather than precise compliance percentages.

The same five-attribute template applies to the operational KPIs in Table 9, whose baselines are the v1 as-is values and whose evidence sources are the KCIs defined in Section 5.3.1.

5.4 Evidence Bases: Policy, Implementation, and Operational Artefacts

To ensure the compliance monitoring framework is auditable and transparent, the CYberSynchrony platform categorizes evidence into three distinct layers: Policy (Governance), Implementation (Configuration), and Operational (Runtime). This structure ensures that the system validates not only that a rule exists on paper (Policy), but also that technical controls are correctly deployed (Implementation) and functioning continuously under real-world conditions (Operational).

The following subsections detail the specific technical tools and project deliverables that serve as the authoritative sources for each category of evidence.

© CYberSynchorny, 2025

5.4.1 Policy and Governance Evidence

The foundational layer of compliance consists of the formal rules, risk assessments, and strategic decisions approved by the organization's management. This evidence is primarily static or semi-static, changing only when organizational strategies are reviewed. The CYBERGOPLUS module monitors this layer by verifying the existence, approval status, and currency of specific governance artifacts generated by the CYberSynchrony toolkits.

- **Cybersecurity Policy Definitions**: The primary source of governance evidence is the Cybersecurity Policy Toolkit (D6.7). The monitoring system verifies the adoption of the 15 core policy tools defined in this deliverable, specifically checking for the "Information Security Strategy" (Tool 1) and "Information Security Governance Policy" (Tool 2). Compliance is measured by the presence of digital signatures or timestamps indicating formal adoption by the Board or CISO.
- **Risk Assessment Records**: Compliance with the "Risk-Driven Governance" principle is evidenced by artifacts from the Risk Scenario Analysis Toolkit (D3.2). The specific evidence required includes the completed "Asset-Based Risk Assessment Workbook" and the "Risk Register". The monitoring system checks the "Last Updated" timestamp of these documents to ensure risk assessments are conducted at least annually as required by the NIS2 methodologies.
- **Data Classification Schemes**: To verify data governance, the system retrieves evidence regarding information labelling protocols. This relies on the Data Protection Policy (Tool 10) from D6.7, which mandates the use of classification schemes such as the Traffic Light Protocol (TLP). The evidence collected is the configuration of data tags within the Secure Data Exchange Platform (D5.1, extended in v2 by the MISP taxonomy tagging described in Section 5.4.3), confirming that data streams are actively tagged (e.g., TLP:RED, TLP:AMBER) before transmission.
- **Team Trustworthiness Assessments**: Complementing the formal risk records above, the trustSense tool provides a team-level risk and readiness assessment that evidences the "Risk-Driven Governance" principle (RM1) from the human and organizational side. Through a structured, consensus-based questionnaire, it produces a Team Trustworthiness score (rated out of 5) across technical, user, and defense roles, together with recommended mitigation actions per category. The monitoring system records the existence and currency of the most recent assessment and its score, and a re-assessment after mitigations provides evidence of improvement over time. Where its recommended mitigations address staff capability, it also contributes to the awareness dimension (TA2). The tool retains no personal data on its users or on assessment outcomes, supporting privacy by design (see Chapter 9).

5.4.2 Implementation and Configuration Evidence

While policy evidence confirms that a rule has been mandated, implementation evidence verifies that the necessary technical controls have been correctly configured and deployed within the IT environment. This layer of the compliance monitoring framework focuses on the "As-Built" state of the infrastructure, auditing configuration files, rule definitions, and system settings to ensure they align with the authorized policies. The CYBERGOPLUS module ingests this evidence directly from the configuration repositories and management consoles of the platform's security tools, including the integration configuration of the additional sources introduced in v2.

- **Identity and Access Configurations**: To verify compliance with the Access Control Policy, the system audits the active configuration of the Keycloak Identity Manager (D6.5). Specific evidence artifacts include the JSON export of the "Realm Settings," which proves that Multi-Factor Authentication (MFA) is set to "Required" for privileged user groups and that password complexity policies (e.g., minimum length, special characters) are technically enforced at the directory level.
- **Endpoint Security Rulesets**: Implementation evidence for asset protection is derived from the configuration of the Wazuh Manager (D4.1). The monitoring system parses the active ossec.conf and agent.conf files to verify that File Integrity Monitoring (FIM) modules are enabled for critical system

© CYberSynchorny, 2025

paths (e.g., /etc/passwd, System32) and that the vulnerability detector module is active with an up-to-date CVE database.

- **Network Detection Signatures**: Compliance with network security standards is evidenced by the active configuration of the Suricata IDS (D4.1). The system audits the suricata.yaml configuration file and the active rule files to confirm that the "Emerging Threats" (ET Open) ruleset is loaded and that the detection engine is configured to inspect traffic on all relevant network interfaces (e.g., eth0, docker0).
- **Infrastructure-as-Code (IaC) Manifests**: For environments deployed via the CYberSynchrony Integration Blueprint (D8.1), compliance is verified through the audit of Kubernetes manifests. The system checks the kube-apiserver.yaml and Ingress Controller configurations (D6.5) to ensure that "Network Policies" are defined to restrict pod-to-pod communication and that TLS certificates are strictly enforced via cert-manager.
- **Security Configuration Assessment Baselines**: Beyond endpoint rulesets, the Wazuh Security Configuration Assessment (SCA) module evidences that each asset is assessed against an authorized hardening baseline. The configuration artifact is the SCA policy assigned to the agent (for example, a CIS benchmark profile), confirming that the asset is continuously checked for secure-configuration conformance. This supports the "Policies and Procedures" control (GOV3) at the configuration layer.
- **Threat-Intelligence Integration Configuration**: To evidence that detection is enriched by curated intelligence, the system audits the configured integration between MISP (D5.3) and the detection tools. The artifact is the MISP REST integration configuration, confirming that indicator feeds and taxonomy tags are pulled into the Wazuh and Suricata pipeline. This establishes, at the configuration layer, the enrichment whose runtime operation is evidenced in Section 5.4.3.

5.4.3 Operational and Runtime Evidence

The final layer of the evidence hierarchy captures the dynamic state of the environment. While configuration evidence proves a control is set to function, operational evidence proves it is actively functioning at a specific point in time. This data is generated continuously by the platform's detection engines, response agents, and user interaction modules. The CYBERGOPLUS module ingests these runtime artifacts to calculate the "Operational Compliance Indicators" (KCIs) defined in Section 5.3.1.

- **Real-Time Threat Detection Logs**: To verify the "Continuous Threat Detection" principle, the system requires proof that monitoring engines are actively inspecting traffic and host activity.
    - **Evidence**: The primary artifact is the stream of Suricata EVE JSON logs from the CYBRITE module (D4.1), which documents every inspected network flow and alert. Wazuh Alert Logs provide evidence of active endpoint monitoring, detailing specific security events such as "Authentication Failure" or "File Integrity Change". In v2, the AI-based IDS (Anomaly Detector, D4.1) additionally contributes anomaly detection records with confidence scores, evidencing behavioral detection of zero-day and denial-of-service activity beyond signature-based rules.
    - **Regulatory Alignment**: This satisfies the NIS2 requirement for "handling incidents" and "security in network and information systems" by proving detection capabilities are operational (EIM1, NS7).
- **Incident Response Execution Records**: Compliance with resilience standards requires evidence that the organization not only detects threats but responds to them according to procedure.
    - **Evidence**: The Automated Response Toolkit (D4.2) generates "Playbook Execution Logs" within the Defensive Cyber Operations (DCO) module. These logs timestamp the exact moment an automated mitigation (e.g., isolating a host) was triggered and completed.
    - **Regulatory Alignment**: This provides the necessary audit trail for "Incident Handling" under NIS2 Article 21 (EIM2).

© CYberSynchorny, 2025

- **Data Integrity and Exchange Verification**: For the "Resilience and Recovery" principle, the system must verify that critical data remains unaltered during operations.
    - **Evidence**: The Secure Data Exchange Platform (D5.1, advanced to v2 in D5.3) and the Evidence Integrity System (D5.2, advanced to v2 in D5.4) produce immutable transaction records anchored in the Hyperledger blockchain. The specific artifact is the "Blockchain Transaction Hash," which serves as cryptographic proof that a specific log entry or shared intelligence file has maintained its integrity from creation to storage.
    - **Regulatory Alignment**: This addresses the GDPR requirements for data integrity and the NIS2 requirements for secure supply chain and data sharing (DS3).
- **Human-Centric Security Metrics**: To validate the effectiveness of the "Cybersecurity Education" objective, the system monitors real-time user behavior and training performance.
    - **Evidence**: The Cybergate (D6.6) platform generates "Mission Outcome Reports" and "Simulation Scores" (e.g., Phishing Defense Success Rate) for every training session. These runtime metrics demonstrate active staff engagement and skill application, moving beyond simple attendance records.
    - **Regulatory Alignment**: This fulfils the "Basic Cyber Hygiene and Training" requirement of NIS2 (TA2).
- **Threat-Intelligence Enrichment and Secure Sharing**: To strengthen the "Continuous Threat Detection" principle and to evidence secure information exchange, the MISP platform (D5.3) operates as the threat-intelligence layer above the detection engines.
    - **Evidence**: At runtime, MISP supplies curated indicator-of-compromise (IoC) events and STIX or IDS-rule exports that are correlated against the Suricata and Wazuh telemetry; the currency of the feed, given by the timestamp of the latest pulled event, evidences that detection is enriched with up-to-date intelligence. For secure sharing, MISP produces PGP or S/MIME signed and encrypted exchange records, sharing-group access logs, and TLP taxonomy tags, which together evidence that intelligence is shared confidentially and with controlled access.
    - **Regulatory Alignment**: The enrichment evidence supports incident detection (EIM1), while the signed, encrypted, and access-controlled sharing supports the GDPR requirement for data integrity and the NIS2 provisions on secure information sharing (DS).
- **External Threat Intelligence and Incident Management**: Complementing the internally generated detection above, the Kitsune platform contributes externally sourced intelligence and a structured incident-management capability.
    - **Evidence**: Kitsune generates credential-leak alerts for the organization's monitored domains, drawn from underground and data-leak sources, alongside curated threat feeds covering ransomware, APT activity, botnet command-and-control, IP reputation, phishing, and malicious hash or URL indicators, delivered over its REST API. Its incident-generation engine produces structured incidents carrying a severity classification, and its incident-response workflow records the handling status, SLA timers, and escalation of each incident. These artifacts evidence both that external threats are detected and that they are triaged and worked through to closure.
    - **Regulatory Alignment**: The detection feeds and credential-leak alerts support incident detection (EIM1), while the structured incidents, severity classification, and SLA-tracked workflow support incident analysis and handling (EIM2) under NIS2 Article 21. Kitsune is integrated as a complementary source that enriches the primary detection and response engines rather than replacing them.

© CYberSynchorny, 2025

The WEKA toolkit does not itself produce runtime compliance evidence; it serves as analytics tooling supporting the development and evaluation of the anomaly-detection models whose outputs appear among the detection evidence above.

To synthesize the methodology detailed in this chapter, Figure 2 presents the overall functional architecture of the compliance monitoring framework. This diagram illustrates how the CYBERGOPLUS module acts as the central processing engine, ingesting raw evidence from the various CYberSynchrony components, validating it against the governance rules defined in D6.3, and ultimately transforming this technical data into the operational indicators, strategic KPIs, and management reports described in the preceding sections.

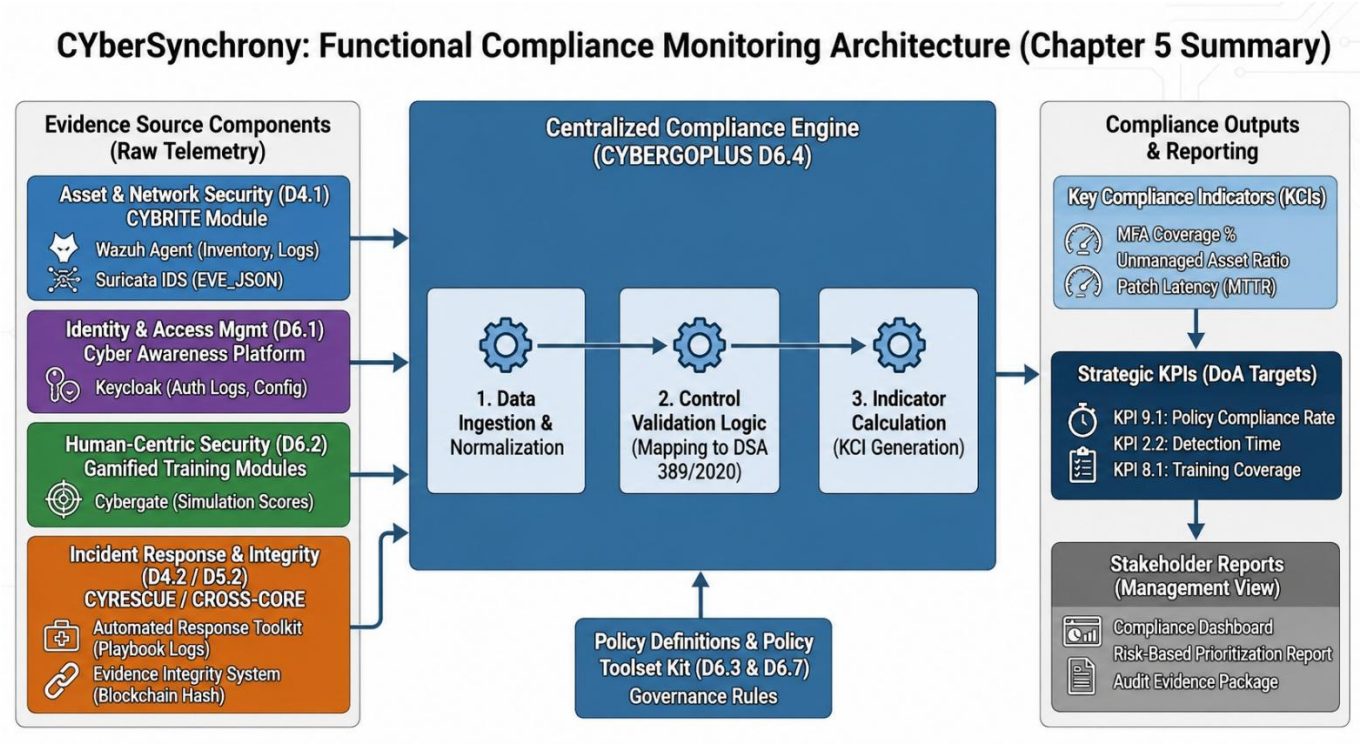

Figure 2: Functional Compliance Monitoring Architecture detailing the flow from raw telemetry to stakeholder reporting.

© CYberSynchorny, 2025

6 System Specifications and Integration Requirements

This chapter translates the high-level compliance methodology into concrete engineering specifications. It defines the functional capabilities the CYBERGOPLUS module must possess, the technical standards it must adhere to, and the integration interfaces required to ingest evidence from the wider CYberSynchrony ecosystem.

6.1 Functional Requirements

The functional specifications for the CYBERGOPLUS module are derived directly from the "Must Have" requirements identified during the End-User Requirements Analysis (D2.2) and the specific governance needs of the NIS2 Directive. The module is designed to fulfil three core functional domains: Policy Governance, Automated Auditing, and Remediation Management.

6.1.1 Policy and Governance Management Functions

To address the end-user requirement for a centralized compliance repository (SQ8), the system must provide digital management for the entire policy lifecycle.

- **REQ-FUNC-01: Policy Registry and Versioning**: The system shall provide a digital registry to store, version, and distribute the 15 Cybersecurity Policy Tools defined in D6.3 (e.g., Information Security Strategy, Access Control Policy). The system must maintain a history of policy revisions, recording the timestamp and the identity of the approving officer (e.g., CISO or Board Member) to satisfy audit trails.
- **REQ-FUNC-02: Role Assignment Matrix**: The system shall provide an interactive matrix to map regulatory roles (e.g., NIS Officer, DPO, Risk Manager) to specific user identities managed in the Keycloak Identity Manager. This function addresses requirement SQ9 (Roles & Responsibilities Matrix) and ensures accountability for specific compliance domains.
- **REQ-FUNC-03: Framework Mapping**: The system shall natively map organizational controls against the specific NIS2 Framework Areas (e.g., Supply Chain Security, Incident Handling) as defined in requirement SQ10. This allows users to view their compliance status filtered by specific NIS2 articles or ISO 27001 domains.

6.1.2 Automated Auditing and Assessment Functions

To reduce the manual workload of gathering audit evidence (SQ1), the system must automate the ingestion and validation of technical telemetry.

- **REQ-FUNC-04: Automated Evidence Ingestion**: The system shall automatically retrieve JSON-formatted log data from the CYBRITE (Wazuh/Suricata) and CROSS-CORE modules at user-defined intervals (e.g., hourly, daily). This function replaces manual log collection for controls such as asset inventory and vulnerability status. In v2, the same ingestion mechanism is extended to the additional evidence sources defined in Section 6.3, including the MISP and Kitsune interfaces.
- **REQ-FUNC-05: Dynamic Asset Reconciliation**: The system shall automatically compare the list of "Active Assets" discovered by Wazuh agents against the "Authorized Inventory" to identify unmanaged devices. This function supports the verification of Asset Management controls (AM1, AM2).
- **REQ-FUNC-06: Vulnerability Risk Mapping**: The system shall ingest CVE reports from Wazuh and SNYK and automatically map them against the organization's defined risk appetite (Low, Medium, High, Critical) to calculate the "Unmitigated Vulnerability" score. This fulfills requirement SQ1 (Internal Regulatory Assessments) by providing real-time risk visibility.

© CYberSynchorny, 2025

6.1.3 Findings Management and Reporting Functions

To satisfy the requirement for management-level reporting (SQ5) and issues management (SQ3), the system must process raw data into actionable insights and track remediation.

- **REQ-FUNC-07: Automated Finding Generation**: The system shall automatically generate a "Non-Compliance Finding" ticket when a Key Compliance Indicator (KCI) falls below the defined threshold (e.g., if "MFA Coverage" drops below 100% for admins). This ensures that deviations are flagged instantly.
- **REQ-FUNC-08: Remediation Workflow Tracking**: The system shall provide a workflow for assigning remediation actions to specific users, tracking due dates, and recording the closure of findings (SQ4). This allows the organization to demonstrate "continuous improvement" to auditors.
- **REQ-FUNC-09: Board-Level Dashboard**: The system shall provide a real-time visualization dashboard displaying the current maturity level (1-5) and compliance percentages. This dashboard must be exportable in standard formats (PDF, CSV) to support the "Management/Board Reporting" requirement (SQ5).

6.2 Technical Specifications

The CYBERGOPLUS module is engineered as a cloud-native application, designed to operate within the Kubernetes-based infrastructure defined in the CYberSynchrony Integration Blueprint. This section details the architectural patterns, technology stack, and security standards required for its implementation.

6.2.1 Architecture Pattern

The system adopts a microservices architecture to ensure modularity and scalability.

- **Containerization**: All application components are packaged as Docker containers, ensuring consistency across development, testing, and production environments.
- **Orchestration**: Deployment is managed via Kubernetes (K8s), which handles service discovery, load balancing, and automated failover. This directly addresses the non-functional requirement for "Automated Failover" (NF4) identified in the end-user analysis.
- **Service Mesh**: Inter-service communication is managed through a service mesh (e.g., Istio or Linkerd) to provide traffic management and observability without modifying application code.
- **API Gateway**: A secure API gateway fronts the module, mediating both the inbound evidence-ingestion connectors and the outbound metrics exposed to the presentation layer (Section 6.3.3). Centralizing ingress at the gateway provides a single point for authentication, rate limiting, and request logging across the growing set of integrated sources.

6.2.2 Technology Stack

The development stack is selected to ensure compatibility with the existing CYberSynchrony components (Wazuh, Keycloak) and to facilitate high-performance data processing.

- **Backend Framework**: The core logic is developed using Python (FastAPI or Django) or Java (Spring Boot). Python is preferred for data processing tasks due to its extensive library support for JSON parsing and API integration, aligning with the Wazuh backend structure.
- **Frontend Framework**: The user interface is built using React or Vue.js. This frontend is designed to be integrated directly into the main CYberSynchrony Platform UI as a micro-frontend, ensuring a unified user experience and "Easy Navigation" (FR14).
- **Integration Libraries**: Evidence connectors are implemented with standard client libraries over the source APIs, including a PyMISP client for the MISP REST interface and REST clients for the Wazuh and

© CYberSynchorny, 2025

Kitsune APIs. Reusing well-supported clients keeps each connector thin and maintainable as sources are added.

6.2.3 Data Storage Strategy

To handle the diverse types of data required for compliance monitoring, the system utilizes a polyglot persistence strategy.

- **Relational Database (SQL)**: PostgreSQL is used to store structured, relational data such as policy definitions, user role assignments, and finding tickets. This ensures transactional integrity for governance records.
- **Log Storage (NoSQL)**: Elasticsearch or OpenSearch is utilized for the storage and rapid querying of high-volume log evidence ingested from Wazuh and Suricata. This aligns with the indexing architecture already deployed in the Monitoring Suite (D4.1).
- **Threat-Intelligence and Indicator Data**: The same document and index store holds the indicator-of-compromise and threat-feed data ingested from MISP and Kitsune, allowing detection telemetry and external intelligence to be correlated in a single query layer. Aggregated KCI and metric values are retained over time to support the trend analysis and maturity computation described in Section 5.3.

6.2.4 Security Standards

The technical implementation adheres to strict security standards to ensure the integrity and confidentiality of compliance data.

- **Transport Security**: All internal (service-to-service) and external (user-to-service) communications must be encrypted using TLS 1.2 or higher. This satisfies the requirement for "Secure Communication Protocols" (NS4).
- **Authentication & Authorization**: The module does not manage user credentials locally. Instead, it relies on OpenID Connect (OIDC) integration with the Keycloak Identity Manager (D6.5) for all authentication requests. Access to specific compliance data is controlled via JWT (JSON Web Token) claims derived from Keycloak roles.
- **Data at Rest and Data Protection**: Stored evidence is encrypted at rest, and evidence categories that contain personal data, such as authentication and training records, are minimized and pseudonymized where the monitoring logic permits, with the full data-protection treatment provided in Chapter 9. Intelligence and evidence exchanged through the Secure Data Exchange Platform (D5.3, v2) additionally inherit its cryptographic protections, including AES-GCM encryption, RSA-OAEP key exchange, and SHA-256 integrity hashing.

6.3 Integration Specifications

To function as the central oversight engine, CYBERGOPLUS requires standardized interfaces to exchange data with the specialized security modules distributed across the CYberSynchrony ecosystem. These integration specifications define the protocols, authentication methods, and data formats necessary to ingest evidence and export compliance status. In this version, the source interfaces are extended to the MISP and Kitsune platforms, and the full data models for all interfaces are provided in Annex I.

6.3.1 Integration with Evidence Sources

The module ingests technical evidence via RESTful APIs handling JSON payloads, aligning with the "Supported Data Formats" requirement (FR20) defined in the end-user requirements. The interface to each source is specified below.

© CYberSynchorny, 2025

- **Input from CYBRITE (Wazuh/Suricata)**: The system interfaces with the Wazuh Manager (D4.1) to retrieve asset, vulnerability, and configuration data, and with the Wazuh Indexer to retrieve detection alerts.
    - **Protocol**: HTTPS / REST API.
    - **Authentication**: JWT obtained from the Wazuh API login endpoint (POST /security/user/authenticate); the token is sent as a Bearer credential on subsequent requests.
    - **Specific Endpoints**:
        - GET /agents: to retrieve the list of active endpoints for inventory reconciliation (AM2).
        - GET /vulnerability/{agent_id}: to retrieve active CVEs for vulnerability compliance (VM1).
        - GET /syscollector/{agent_id}/...: to retrieve hardware and OS details for configuration auditing.
        - GET /sca/{agent_id}: to retrieve Security Configuration Assessment results for hardening conformance (GOV3).
        - Wazuh Indexer query on the wazuh-alerts-* index: to retrieve endpoint alerts and the fused Suricata EVE records for detection evidence (EIM1, NS7).
    - **Data Format**: JSON output containing asset IDs, rule identifiers and levels, timestamps, and event types.
- **Input from CROSS-CORE (Evidence Integrity)**: The system interfaces with the Evidence Integrity System (D5.2, advanced to v2 in D5.4) to verify the immutability of logs.
    - **Protocol**: HTTPS / REST API.
    - **Authentication**: Keycloak-issued OIDC JWT obtained from POST `/realms/CYberSynchrony/protocol/openid-connect/token`, sent as a Bearer credential.
    - **Specific Endpoints**: POST `/api/crosscore/verify/result` — re-hashes the evidence artifact and checks it against the anchored `resultHash`.
    - **Data Format**: JSON. Response: `{ match: boolean, anchoredTxId: string }`.
    - **Function**: Log Integrity Verification. The compliance module queries the Hyperledger ledger to confirm that the hash of a specific evidence log matches the stored immutable record (DS3).
- **Input from MISP (Threat Intelligence)**: The system interfaces with the MISP platform (D5.3) to retrieve curated threat intelligence for detection enrichment and to record secure-sharing evidence.
    - **Protocol**: HTTPS / REST API.
    - **Authentication**: A MISP automation API key transmitted in the Authorization header.
    - **Specific Endpoints**:
        - POST /events/restSearch: to retrieve indicator-of-compromise events, with the query constrained by timestamp, tag, and TLP level.
        - GET /attributes/restSearch: to retrieve individual indicators (IPs, hashes, URLs) for correlation against telemetry.
        - GET /feeds: to confirm the configured intelligence feeds and their last-fetched status.

© CYberSynchorny, 2025

- **Data Format**: JSON events and attributes carrying TLP taxonomy tags; the same search endpoints can return the intelligence in STIX 2 or in Suricata and Snort rule format via the returnFormat parameter, for direct consumption by the detection engines.

- **Input from Kitsune (External Threat Intelligence)**: The system interfaces with the Kitsune platform to retrieve externally sourced intelligence, credential-leak alerts, and structured incidents.
  - **Protocol**: HTTPS / REST API, with the service base at https://home.kitsuneapp.eu.
  - **Authentication**: Bearer token, obtained through the login endpoint and maintained via the refresh-token endpoint.
  - **Specific Endpoints**:
    - POST /api/do/v2/lookup: to search across the underground and leak datasets for a given entity.
    - POST /api/do/entity/{type}: typed entity lookups, where {type} is one of email-address, email-domain, ip, bank, card, or crypto, used for credential-leak and exposure checks against the organization's monitored identifiers.
    - POST /api/i/bots and GET /api/i/bots/{indexfile}: to retrieve stealer-log and bot records.
    - The leak, threat-feed, and alert endpoints: to retrieve credential-leak findings, curated feeds (ransomware, APT, botnet command-and-control, IP reputation, phishing, and malicious hash or URL indicators), and the structured incidents with severity classification.
  - **Data Format**: JSON alerts and incidents carrying severity, tags, a matched-string field, and timestamps, together with credential-leak records.

6.3.2 Integration with Identity Services

The system integrates with the Keycloak Identity Manager (D6.5) both as its own authentication provider and as a source of identity and access evidence. The compliance module does not store or manage user credentials; it federates authentication to Keycloak and queries Keycloak for the evidence required to verify the identity controls.

- **Service**: Keycloak Identity Manager (D6.5).
- **Integration Type**: OpenID Connect (OIDC) for authentication and token validation, with SAML supported where the host environment requires it. Administrative evidence is retrieved through the Keycloak Admin REST API.
- **Module Authentication**: The module is registered as a confidential OIDC client. Interactive user sessions use the Authorization Code flow, while the module's service-to-service calls to the Admin REST API use the Client Credentials grant with a dedicated service account holding the realm-management roles view-users and view-events.
- **Specific Endpoints**:
  - GET /realms/{realm}/.well-known/openid-configuration: discovery of the realm's OIDC endpoints.
  - POST /realms/{realm}/protocol/openid-connect/token: token issuance for user sessions and the Client Credentials grant for service calls.
  - GET /realms/{realm}/protocol/openid-connect/certs: the JWKS used to verify token signatures locally.

© CYberSynchorny, 2025

- GET /admin/realms/{realm}/users and GET /admin/realms/{realm}/users/{id}/credentials: to determine the user population and whether an OTP credential is configured for each user, evidencing MFA enrollment.
- GET /admin/realms/{realm}/authentication/flows and GET /admin/realms/{realm}/authentication/required-actions: to verify that the browser authentication flow enforces OTP for privileged groups (configuration evidence for IAM4).
- GET /admin/realms/{realm}/events?type=LOGIN: to retrieve login events as runtime evidence of authentication activity.

- **Token Claims Inspected**:
  - acr (Authentication Context Class Reference) and amr (Authentication Methods References): to confirm that a session was established with strong, multi-factor authentication (IAM4).
  - realm_access.roles and resource_access.{client}.roles: to map Keycloak roles (e.g., cis_admin, auditor) to the permissions defined in the CYBERGOPLUS Role Assignment Matrix (REQ-FUNC-02), supporting role-based access control (IAM1).
  - sub and preferred_username: to associate evidence and findings with an accountable identity without storing credentials.
- **Data Flow and Control Mapping**:
  - **MFA Coverage (IAM4)**: combining the configuration evidence (OTP required in the browser flow for privileged groups), the per-user credential check, and the acr and amr claims observed on sessions, the module computes the proportion of privileged identities and logins protected by multi-factor authentication.
  - **Role Assignment (IAM1)**: the module reconciles the roles asserted by Keycloak against the Role Assignment Matrix, flagging any identity whose roles fall outside the authorized mapping.

6.3.3 Integration with Presentation Layer

The compliance module surfaces its results to the unified platform interface through a read-only API exposed by the API gateway introduced in Section 6.2.1.

- **Target**: CYberSynchrony Platform UI (D8.1), into which the module is embedded as a micro-frontend.
- **Mechanism**: CYBERGOPLUS exposes a secure, read-only REST API through its API gateway, allowing the main dashboard to fetch real-time KCI and KPI metrics (e.g., "Policy Compliance Rate") and the current maturity posture for visualization in the unified view.
- **Authentication**: Requests carry a Keycloak-issued JWT; the gateway validates the signature against the realm JWKS and scopes each response to the roles asserted in the token, so that a user sees only the compliance data their role permits (consistent with Section 6.3.2).
- **Specific Endpoints (exposed by CYBERGOPLUS)**:
  - GET /api/v1/compliance/status: returns the ComplianceStatusJSON object for the dashboard summary view.
  - GET /api/v1/compliance/findings: returns the list of open non-compliance findings, filterable by control, severity, or assessment area.
  - GET /api/v1/compliance/kpis: returns the current KPI scores together with their methodology metadata (baseline, target, measurement window) as defined in Section 5.3.4.
  - GET /api/v1/compliance/report?format={pdf|csv}: returns the exportable board-level report required by REQ-FUNC-09.

© CYberSynchorny, 2025

- **Data Object (ComplianceStatusJSON)**: a JSON document containing:
  - overall_maturity_level (1 to 5) and overall_compliance_percentage;
  - assessment_areas: an array of objects, each with the area name and its maturity_level and compliance_percentage;
  - kpis: an array of objects, each with the KPI identifier, current value, baseline, target, and measurement_window;
  - open_findings: an array of objects, each with the finding identifier, the affected control, severity, asset_id, status, and due_date;
  - generated_at: the timestamp of the snapshot.
- **Protocol and Format**: HTTPS / REST, returning JSON.

The required data models above are described in full in Annex I.

The diagram below indicates the overall architecture of CYBERGOPLUS:

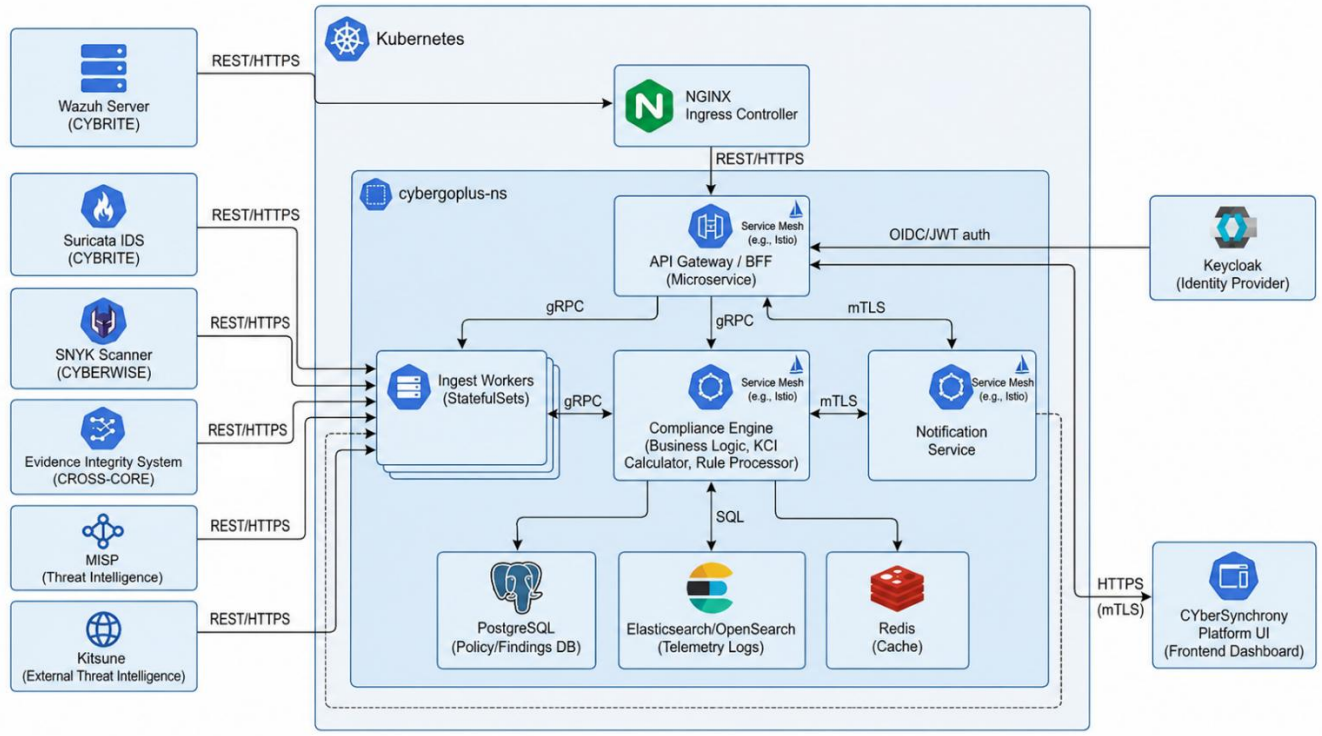

Figure 3: CYBERGOPLUS Technical Architecture. This diagram illustrates the module's microservices deployment within Kubernetes, detailing data ingestion from external sources (CYBRITE, CROSS-CORE, MISP, and Kitsune), internal processing via the Compliance Engine

7 End-to-End Compliance Workflows

This chapter presents the compliance monitoring framework as a set of concrete, end-to-end workflows. Where the preceding chapters defined the evidence sources, indicators, and integration interfaces as discrete elements, this chapter shows them operating together. Each workflow traces a single thread of evidence as it moves across multiple modules, from its generation in an integrated component, through ingestion and validation in the CYBERGOPLUS Compliance Engine, to the resulting compliance indicator, finding, and remediation action.

The workflows sequence elements already specified in this deliverable into the operational paths an organization follows in practice, drawing on the evidence sources of Chapter 5, the KCIs and measurement logic of Section 5.3, and the integration interfaces of Chapter 6. In line with the project timeline, they describe the designed and configured workflows of the monitoring system rather than retrospective records of pilot execution.

Each workflow is presented in a common structure: the triggering condition, the evidence generated and the component that produces it, the path through CYBERGOPLUS (ingestion, KCI evaluation, and finding generation), the remediation and re-validation step, and the controls and KPIs the workflow evidences.

The chapter presents the following workflows:

- **The Compliance Monitoring Loop** (Section 7.1): the generic end-to-end cycle common to every workflow, from evidence to finding to re-validation and maturity update.
- **Phishing Incident Detection and Response** (Section 7.2): a cross-module workflow spanning detection, threat-intelligence enrichment, automated containment, identity remediation, and staff retraining.
- **Vulnerability Detection and Remediation** (Section 7.3): the path from CVE and code-scan findings through risk mapping to remediation tracking and the remediation-time KPI.
- **Identity and Access Compliance Monitoring** (Section 7.4): the path from Keycloak configuration and login evidence to the MFA-coverage indicator and any resulting finding.

7.1 The Compliance Monitoring Loop

Every workflow in this chapter is an instance of a single underlying cycle. The CYBERGOPLUS module runs this loop continuously, combining the outputs of the integrated components into a repeatable path from raw evidence to a verified compliance outcome. The stages are as follows.

- **Evidence Generation**: The integrated components continuously produce runtime artifacts: detection telemetry from Wazuh, Suricata, and the AI-based IDS; identity events from Keycloak; vulnerability findings from Wazuh and SNYK; threat intelligence from MISP and Kitsune; and response logs from the DCO playbooks.
- **Ingestion**: At the intervals defined in REQ-FUNC-04, the CYBERGOPLUS connectors retrieve these artifacts through the interfaces specified in Chapter 6, normalizing each into a common evidence record.
- **Correlation**: Externally sourced intelligence from MISP and Kitsune is correlated against the internal detection telemetry, so that indicators and observed activity are evaluated together rather than in isolation.
- **KCI Evaluation**: Each evidence record is evaluated against the atomic KCI logic of Section 5.3.1, producing a Pass or Fail flag with a timestamp and an Asset ID, tied to a specific DSA 389/2020 control.
- **Aggregation**: KCIs are aggregated into the metrics and KPIs of Section 5.3.2, the methodology of Section 5.3.4 where relevant, and the maturity levels of Section 5.3.3.
- **Finding Generation**: Where a KCI falls below its threshold, the engine raises a Non-Compliance Finding (REQ-FUNC-07), recording the control, the affected asset, and the severity.

© CYberSynchorny, 2025

- **Remediation**: The finding is assigned, tracked, and closed through the remediation workflow (REQ-FUNC-08). Depending on the control, remediation may be a manual action or an automated response executed by a DCO playbook.
- **Re-validation**: The next evidence cycle re-evaluates the same KCI. A subsequent Pass confirms the control has been restored and closes the finding; a continued Fail keeps it open.
- **Reporting**: The updated indicators, KPIs, and maturity levels are exposed to the board-level dashboard (REQ-FUNC-09) through the presentation API of Section 6.3.3.

The workflows in Sections 7.2 to 7.4 instantiate this loop for specific scenarios. The remainder of this section traces the loop once at the level of the actual API calls.

7.1.1 The Loop in API Terms

To make the loop concrete, this section traces a single control, Asset Management (AM2), from end to end at the level of the specific API calls. The workflow combines the Suricata flow telemetry, the Wazuh Indexer, and the Wazuh Manager, and corresponds to the Dynamic Asset Reconciliation function (REQ-FUNC-05).

- **Evidence generation**: Suricata records network flows as EVE JSON (event_type: flow), which are forwarded to the Wazuh Indexer, while Wazuh maintains the registered asset inventory through its agents and the syscollector module.
- **Ingestion and retrieval**: The connector authenticates to the Wazuh Manager with POST /security/user/authenticate, receiving a JWT. It retrieves the authorized inventory with GET /agents?status=active and the registered addresses with GET /syscollector/{agent_id}/netaddr. It then queries the Wazuh Indexer over the wazuh-alerts-* index for the distinct source and destination IP addresses observed in Suricata flow events within the measurement window.
- **KCI evaluation**: For each active IP observed by Suricata, the engine checks membership in the authorized inventory and applies the rule "an active IP observed but absent from the authorized inventory implies an unmanaged asset", emitting a Pass or Fail per IP with a timestamp and the IP as the Asset ID. This KCI verifies control AM2 (Inventory of Assets).
- **Aggregation**: The engine computes the Unmanaged Asset Ratio (unmanaged IPs over total active IPs) defined in Table 9, contributing to the Policy Compliance Rate (KPI 9.1) and to the Asset Security maturity area.
- **Finding generation**: If the ratio exceeds the configured threshold, the engine raises a Non-Compliance Finding (REQ-FUNC-07) recording control AM2, the offending IP addresses, and a severity.
- **Remediation**: The finding is assigned through the remediation workflow (REQ-FUNC-08); the asset is either brought under management, by deploying a Wazuh agent so that it subsequently appears in GET /agents, or formally decommissioned.
- **Re-validation**: On the next cycle, the same query and KCI are re-run. Once the IP is present in the authorized inventory or no longer observed, the KCI returns Pass and the finding is closed.
- **Reporting**: The recomputed Asset Security maturity and the AM2 indicator are surfaced to the dashboard via GET /api/v1/compliance/status.

Following the submission of this blueprint, Task 6.7 shifts focus to the active engineering of the defined modules. The immediate next steps involve transforming the architectural specifications into working software components.

Primary activities for the next period include:

- **Environment Setup:** Establishing the dedicated cybergoplus-ns namespace within the development Kubernetes cluster and configuring base services (PostgreSQL, Kafka).

© CYberSynchorny, 2025

- **Connector Development:** Coding the southbound adapters required to interface with the Wazuh REST API and Keycloak OIDC endpoints.
- **Engine Logic Implementation:** Developing the core business logic in Python to process incoming JSON telemetry and calculate the operational KCIs defined in Chapter 5.
- **Alpha Testing:** Conducting initial integration tests with technical partners (BYTE, IANUS) to validate data ingestion flows before pilot deployment. Capacity-Building Needs

7.2 Phishing Incident Detection and Response

Phishing is the most cross-cutting incident type in the pilot environments, and its handling exercises the largest number of components in a single sequence. This workflow follows a phishing incident from the initial SIEM trigger through enrichment, containment, identity remediation, and staff retraining, and shows the compliance evidence each step produces. The response procedure is the phishing playbook shown in Figure 4; the description below maps each step to the component that performs it and the evidence it generates.

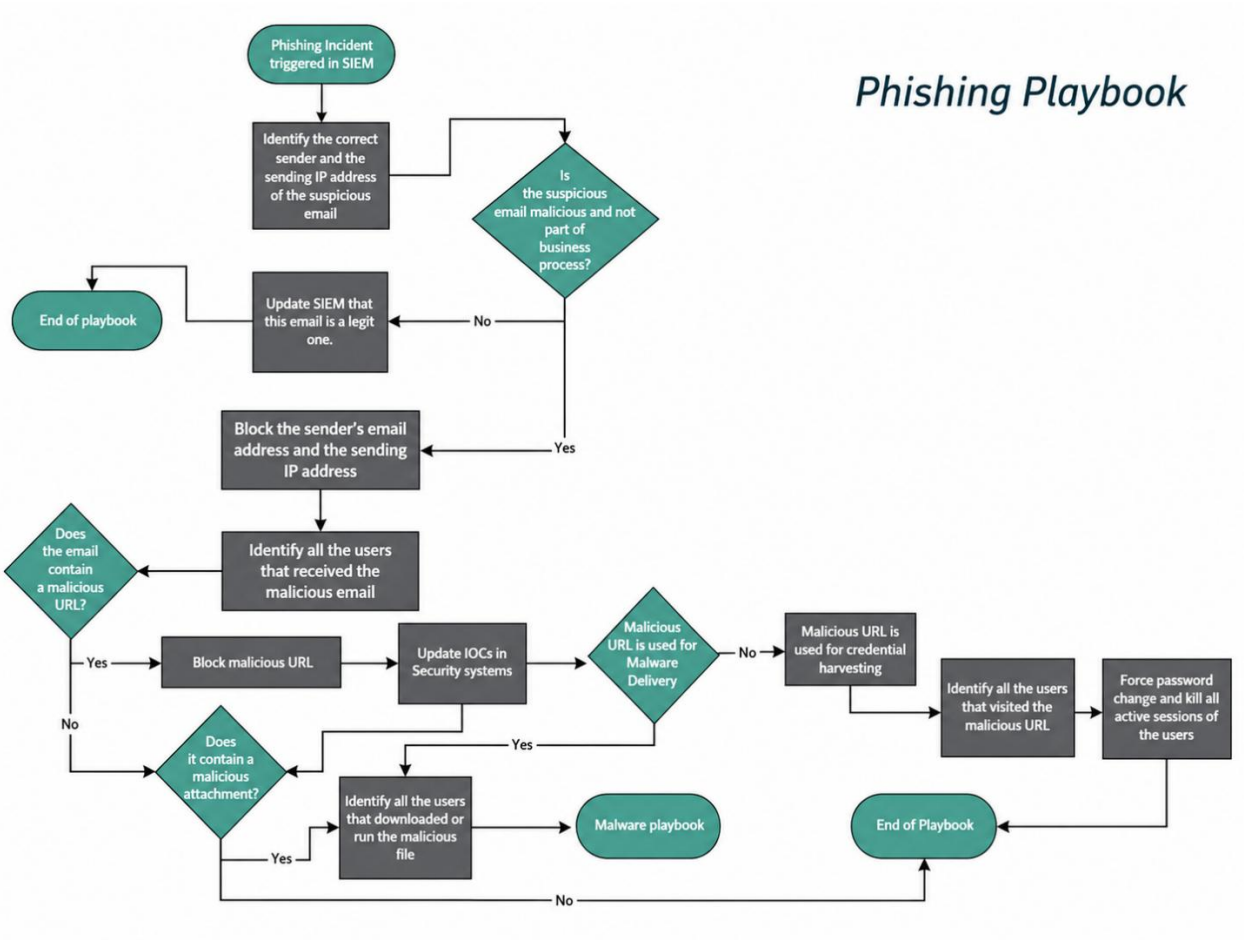

Figure 4: Phishing Incident Response Playbook.

1. **Trigger**: A phishing incident is raised when Wazuh correlates suspicious email and endpoint activity into an alert in the wazuh-alerts index.
2. **Identification**: The sending address and source IP of the suspicious email are identified from the Wazuh log analysis.
3. **Maliciousness Determination**: The sender IP, any embedded URLs, and attachment hashes are checked against the threat intelligence, through MISP indicator lookups (GET /attributes/restSearch) and the

Kitsune malicious IP, URL, and hash feeds. If the indicators are benign and the message belongs to a legitimate business process, the SIEM is updated and the playbook ends; otherwise the response proceeds.

4. **Containment**: The DCO Automated Response Toolkit (D4.2) blocks the sender address and the source IP, recording the action in a Playbook Execution Log.
5. **Scope Identification**: All users who received the message are identified from the email and endpoint logs.
6. **URL and Attachment Handling**: If a malicious URL is present, it is blocked and the new indicators are written back into the security systems (recorded in MISP via POST /events and propagated to the detection rules). The workflow then branches: a malware-delivery URL or a malicious attachment routes the affected users to the malware playbook, while a credential-harvesting URL triggers identity remediation.
7. **Identity Remediation**: For each user who visited a credential-harvesting URL, the workflow forces a password change and terminates active sessions through Keycloak (POST /admin/realms/{realm}/users/{id}/logout, with an UPDATE_PASSWORD required action set on the account), restoring the integrity of the affected identities.
8. **Retraining**: The affected users are enrolled in a targeted Cybergate (D6.6) training module, which generates a mission outcome record on completion.

Across this single sequence, the workflow combines six components (Wazuh, MISP, Kitsune, the DCO playbooks, Keycloak, and Cybergate) and produces the following compliance evidence:

- **Detection and enrichment**: the SIEM trigger and the intelligence lookups evidence Continuous Threat Detection (EIM1, NS7) and contribute to the Detection Latency metric (KPI 2.2).
- **Containment and handling**: the playbook execution logs evidence Incident Analysis and Handling (EIM2) under NIS2 Article 21 and contribute to the Resolution Rate metric (KPI 2.4).
- **Identity remediation**: the forced password reset and session termination evidence Strong Authentication and identity integrity (IAM4).
- **Intelligence sharing**: the indicator write-back to MISP evidences secure threat-information sharing (DS).
- **Retraining**: the Cybergate mission record evidences Awareness Training (TA2) and contributes to Training Coverage (KPI 8.1).

7.3 Vulnerability Detection and Remediation

Vulnerability Detection and Remediation

This workflow tracks a vulnerability from its detection through risk-based prioritization to remediation and re-validation, combining the host-level findings of Wazuh with the code, dependency, and container findings of SNYK. It is the path that produces the organization's vulnerability-status posture and its remediation-time KPI.

1. **Trigger**: A vulnerability is recorded either by the Wazuh Vulnerability Detector, which correlates installed packages against its CVE database to flag a CVE on a monitored host, or by SNYK, which flags a finding in application code, a dependency, a container image, or an infrastructure-as-code manifest.
2. **Ingestion**: At the interval defined in REQ-FUNC-04, the connectors retrieve the findings: from the Wazuh Manager with GET /vulnerability/{agent_id} (authenticated with a JWT from POST /security/user/authenticate), and from the SNYK API issues endpoint (for example, GET

© CYberSynchorny, 2025

/rest/orgs/{org_id}/issues, authenticated with a SNYK API token). Each finding carries a severity, an identifier (CVE or CWE), and a location.

3. **Risk Mapping**: The engine maps each finding against the organization's defined risk appetite (Low, Medium, High, Critical) to compute the Unmitigated Vulnerability score (REQ-FUNC-06), so that remediation is prioritized by risk rather than by raw count.
4. **KCI Evaluation**: The atomic KCI applies the rule "no open Critical or High vulnerability remains on the asset or project", emitting a Pass or Fail with a timestamp and the Asset ID. This KCI verifies control VM1 (Vulnerability Scanning).
5. **Aggregation**: For each remediated finding, the engine records the Mean Time to Remediate (the interval between the detection timestamp and the patch or fix timestamp) defined in Table 9, contributing to the Platform Vulnerability Remediation Time KPI (KPI 1.5 and KPI 3.5) and to the Vulnerability Status maturity area.
6. **Finding Generation**: Where an open Critical or High vulnerability is present, the engine raises a Non-Compliance Finding (REQ-FUNC-07) recording control VM1, the affected asset or project, the CVE or CWE, and the severity.
7. **Remediation**: The finding is assigned through the remediation workflow (REQ-FUNC-08); the host is patched, the dependency upgraded, the container rebuilt, or the manifest corrected, and the action is tracked to closure.
8. **Re-validation**: On the next scan cycle, the same retrieval and KCI are re-run. Once the CVE or finding no longer appears, the KCI returns Pass, the finding is closed, and the remediation time is finalized.

This workflow evidences Vulnerability Scanning (VM1), its remediation timestamps contribute to the Vulnerability Remediation Time KPI (KPI 1.5 and KPI 3.5), and its aggregate pass rate updates the Vulnerability Status maturity area.

7.4 Identity and Access Compliance Monitoring

This workflow verifies the organization's identity controls on a continuous basis, drawing entirely on the Keycloak Identity Manager and the integration detailed in Section 6.3.2. Unlike the incident-driven workflows above, it runs as a scheduled evaluation and on any change to the realm configuration, and it produces both the multi-factor-authentication posture and the role-conformance posture.

1. **Trigger**: The workflow runs on the scheduled evaluation interval (REQ-FUNC-04) and whenever a relevant change is made to the Keycloak realm, such as a new user, a credential change, or a role assignment.
2. **Ingestion**: The connector authenticates as its service account (POST /realms/{realm}/protocol/openid-connect/token, client_credentials grant) and retrieves the user population with GET /admin/realms/{realm}/users, the per-user credential set with GET /admin/realms/{realm}/users/{id}/credentials to determine OTP enrollment, the browser flow and required actions with GET /admin/realms/{realm}/authentication/flows and GET /admin/realms/{realm}/authentication/required-actions, and recent login activity with GET /admin/realms/{realm}/events?type=LOGIN. The acr, amr, and realm_access.roles claims are read from observed session tokens.
3. **KCI Evaluation (MFA)**: The atomic KCI applies the rule "MFA is required in the browser flow for privileged groups, each privileged identity has an OTP credential, and observed privileged logins carry an acr or amr value indicating multi-factor authentication", emitting a Pass or Fail per identity. This verifies control IAM4 (Strong Authentication).

© CYberSynchorny, 2025

4. **KCI Evaluation (RBAC)**: A second KCI applies the rule "each identity's Keycloak roles fall within the authorized Role Assignment Matrix", emitting a Pass or Fail per identity. This verifies control IAM1 (Role-Based Access Control).
5. **Aggregation**: The engine computes MFA Coverage (the proportion of privileged identities and logins protected by multi-factor authentication), contributing to the Policy Compliance Rate (KPI 9.1) and the Identity Control maturity area, alongside the proportion of identities whose roles conform to the matrix.
6. **Finding Generation**: Where MFA coverage falls below its threshold (for example, below 100% for administrators) or an identity holds roles outside the authorized mapping, the engine raises a Non-Compliance Finding (REQ-FUNC-07) recording the control, the affected identity, and the severity.
7. **Remediation**: The finding is assigned through the remediation workflow (REQ-FUNC-08). MFA gaps are closed by setting the CONFIGURE_TOTP required action on the account (PUT /admin/realms/{realm}/users/{id}), and role deviations are corrected through the realm role-mapping endpoints; both are tracked to closure.
8. **Re-validation**: On the next cycle, the same queries and KCIs are re-run. Once the identity has enrolled MFA or its roles match the matrix, the KCI returns Pass and the finding is closed.

This workflow evidences Strong Authentication (IAM4), feeding MFA Coverage within the Policy Compliance Rate (KPI 9.1), and Role-Based Access Control (IAM1), with the aggregate pass rate updating the Identity Control maturity area.

7.5 Governance and Risk Verification

Not every control is verified from runtime telemetry. The governance and risk controls are verified from the governance artifacts themselves, on a periodic review cycle. This workflow checks the currency and completeness of the policy, role, and risk evidence, combining the CYBERGOPLUS Policy Registry and Role Assignment Matrix with the D3.2 risk records and the trustSense team assessment.

1. **Trigger**: The workflow runs on the governance review interval (REQ-FUNC-04, set to the organization's review cadence) and whenever a governance artifact is added or revised.
2. **Ingestion**: The engine reads its own Policy Registry (REQ-FUNC-01) and Role Assignment Matrix (REQ-FUNC-02) from the governance store, reconciles the matrix against the Keycloak roles (GET /admin/realms/{realm}/roles), and retrieves the currency of the external governance artifacts: the Last Updated timestamps of the Risk Register and the Asset-Based Risk Assessment Workbook (D3.2), and the date and score of the most recent trustSense assessment.
3. **KCI Evaluation (Policies)**: The atomic KCI applies the rule "each of the 15 policy tools has a current, formally adopted digital record, signed or timestamped and reviewed within the review period", emitting a Pass or Fail per policy. This verifies control GOV3 (Policies and Procedures).
4. **KCI Evaluation (Roles)**: A second KCI applies the rule "each required regulatory role, such as the NIS Officer, DPO, and Risk Manager, is assigned to an accountable identity in the Role Assignment Matrix", emitting a Pass or Fail per role. This verifies control GOV1 (Roles and Responsibilities).
5. **KCI Evaluation (Risk)**: A third KCI applies the rule "the risk assessment has been updated within the required interval, at least annually, and a current trustSense team assessment exists", emitting a Pass or Fail. This verifies control RM1 (Risk Management Methodology).
6. **Aggregation**: The engine computes the Policy Adoption Rate (the proportion of the 15 tools that are current), the role-assignment completeness, and the risk-assessment currency, contributing to the Policy Compliance Rate (KPI 9.1), the Risk Monitoring Services count (KPI 5.1), and the Governance and Strategy maturity area.

© CYberSynchorny, 2025

7. **Finding Generation**: Where a policy is missing or stale, a required role is unassigned, or a risk assessment is overdue, the engine raises a Non-Compliance Finding (REQ-FUNC-07) recording the control and the affected artifact.
8. **Remediation**: The finding is assigned through the remediation workflow (REQ-FUNC-08); the policy is adopted or reviewed, the role is assigned, or the risk or trustSense assessment is conducted, and the action is tracked to closure.
9. **Re-validation**: On the next review cycle, the same checks are re-run. Once the artifact is current and complete, the KCI returns Pass and the finding is closed.

This workflow evidences Policies and Procedures (GOV3), feeding the Policy Adoption Rate within the Policy Compliance Rate (KPI 9.1), Roles and Responsibilities (GOV1), and Risk Management Methodology (RM1), which contributes to the Risk Monitoring Services count (KPI 5.1), with the aggregate pass rate updating the Governance and Strategy maturity area.

7.6 Resilience and Integrity Verification

This workflow verifies that the organization can withstand and recover from disruption and that its critical data remains unaltered. It combines the backup and continuity records with the cryptographic integrity proofs produced by the Evidence Integrity System, verified through the CROSS-CORE interface of Section 6.3.1.

1. **Trigger**: The workflow runs on the resilience review interval (REQ-FUNC-04) and after each scheduled backup or continuity exercise.
2. **Ingestion**: The engine ingests the backup completion records from the backup routines and the business continuity drill records from the continuity register, and it queries the Evidence Integrity System through the CROSS-CORE interface (RESTAPI over HTTPS) to obtain the Blockchain Transaction Hash for each backup archive and critical log, confirming whether the stored hash matches the current artifact.
3. **KCI Evaluation (Backup)**: The atomic KCI applies the rule "the scheduled backup completed within its window and the archive's integrity hash is valid and unaltered", emitting a Pass or Fail per backup. This verifies control DS3 (Backup and Recovery).
4. **KCI Evaluation (Continuity)**: A second KCI applies the rule "a business continuity drill was executed within the required interval and recorded", emitting a Pass or Fail. This verifies control BCR2 (Business Continuity).
5. **KCI Evaluation (Integrity)**: A third KCI applies the rule "critical log entries and shared files carry valid, unaltered blockchain hashes", emitting a Pass or Fail per artifact. This verifies control DS (Data Integrity).
6. **Aggregation**: The engine computes the Integrity Verification Rate (the proportion of backup archives and logs with valid hashes) defined in Table 9, alongside the backup success rate and the continuity-drill currency, contributing to the Risk Monitoring Services count (KPI 5.1) and the Resilience maturity area.
7. **Finding Generation**: Where a backup failed or is missing, a continuity drill is overdue, or an integrity hash does not match, the engine raises a Non-Compliance Finding (REQ-FUNC-07) recording the control and the affected artifact, with an integrity mismatch raised at high severity.
8. **Remediation**: The finding is assigned through the remediation workflow (REQ-FUNC-08); the backup is re-run, the continuity drill is scheduled and conducted, or the integrity mismatch is investigated, and the action is tracked to closure.
9. **Re-validation**: On the next cycle, the same checks are re-run. Once the backup completes with a valid hash, the drill is current, and the integrity proofs match, the KCI returns Pass and the finding is closed.

© CYberSynchorny, 2025

This workflow evidences Backup and Recovery (DS3), Business Continuity (BCR2), and Data Integrity (DS), feeding the Integrity Verification Rate; all three contribute to the Risk Monitoring Services count (KPI 5.1) and update the Resilience maturity area.

7.7 Control Coverage Summary

The workflows above, together with the continuous monitoring loop, exercise every control in the framework. The following matrix maps each DSA 389/2020 control to the workflow that verifies it and its principal evidence source, confirming that no control is left unmonitored.

Table 11. Controls and Verification Workflows

Control	Control Name	Verifying Workflow	Principal Evidence Source
GOV1	Roles and Responsibilities	7.5 Governance and Risk Verification	CYBERGOPLUS Role Assignment Matrix
GOV3	Policies and Procedures	7.5 Governance and Risk Verification	CYBERGOPLUS Policy Registry (D6.3)
RM1	Risk Management Methodology	7.5 Governance and Risk Verification	D3.2 Risk Register; trustSense
AM1	Asset Management	7.1.1 Asset Reconciliation	Wazuh inventory
AM2	Inventory of Assets	7.1.1 Asset Reconciliation	Wazuh inventory; Suricata flows
IAM1	Role-Based Access Control	7.4 Identity and Access Compliance	Keycloak roles
IAM4	Strong Authentication	7.4 Identity and Access Compliance (and 7.2)	Keycloak MFA configuration and events
NS4	Secure Communication Protocols	7.1 Monitoring Loop (continuous KCI)	Suricata TLS records
NS7	Intrusion Detection	7.2 Phishing Detection and Response	Suricata and Wazuh alerts
VM1	Vulnerability Scanning	7.3 Vulnerability Detection and Remediation	Wazuh and SNYK findings
EIM1	Incident Detection	7.2 Phishing Detection and Response	Detection telemetry; MISP and Kitsune
EIM2	Incident Analysis	7.2 Phishing Detection and Response	DCO playbook logs; Kitsune incidents

© CYberSynchorny, 2025

BCR2	Business Continuity	7.6 Resilience and Integrity Verification	Continuity drill records
DS3	Backup and Recovery	7.6 Resilience and Integrity Verification	Backup records; Evidence Integrity hashes
DS	Data Integrity	7.6 Resilience and Integrity Verification (and 7.2)	Evidence Integrity blockchain hashes; MISP signed sharing
TA2	Awareness Training	7.2 Phishing Detection and Response	Cybergate mission records
GOV1	Roles and Responsibilities	7.5 Governance and Risk Verification	CYBERGOPLUS Role Assignment Matrix
GOV3	Policies and Procedures	7.5 Governance and Risk Verification	CYBERGOPLUS Policy Registry (D6.3)

© CYberSynchorny, 2025

8 Per-Pilot Integrated Narratives

The preceding chapters described the compliance monitoring framework as a general capability. This chapter applies it to the three sectoral pilots, telling each use case's story end-to-end in a single place. The sector-specific requirements were introduced in Section 4.2 and the workflows in Chapter 7; this chapter brings them together per pilot, so that the configuration, the workflows exercised, the resulting compliance posture, and the operational considerations for each use case can be read as one coherent narrative.

The three pilots represent distinct critical sectors with distinct compliance profiles: the Apostolos Loucas Medical Center (UC1, Healthcare, supported by eBOS), the EGL public administration pilot (UC2, Public Administration, supported by MAESTRO), and the UPRC academic pilot (UC3).

In line with the project timeline, each narrative describes the designed and configured application of the framework to the pilot, including the components deployed and the controls in scope, rather than reporting longitudinal measured results. Where an indicative figure would otherwise appear, the narrative states the measurement approach and the evidence source instead, consistent with the methodology of Chapter 5.

Each pilot narrative follows a common structure:

- **Pilot Profile and Sector Context**: the organization, its critical assets, and the regulatory drivers specific to its sector.
- **Configuration**: the components and evidence sources deployed for the pilot, and the controls in scope.
- **Workflows Exercised**: the Chapter 7 workflows most relevant to the pilot, and how they apply to its assets.
- **Compliance Posture**: the assessment areas and the maturity the framework produces for the pilot, expressed as the measurement approach.
- **Operational Considerations**: the sector-specific constraints, risks, and findings that shape monitoring in the pilot.

The chapter is organized as follows:

- **Section 8.1**: UC1, Healthcare (Apostolos Loucas Medical Center).
- **Section 8.2**: UC2, Public Administration (EGL).
- **Section 8.3**: UC3, Academia (UPRC).
- **Section 8.4**: Cross-Pilot Observations, drawing the three narratives together to show the harmonized framework applied across different sectors, in line with the Task 6.6 harmonization objective.

8.1 UC1: Healthcare (Apostolos Loucas Medical Center)

Pilot profile and sector context. Apostolos Loucas Medical Center (AMC) is the largest and most comprehensive outpatient medical center in Cyprus, and it leads the healthcare use case. Its environment combines clinical, administrative, and imaging systems across two buildings (Block A and Block B), with the internal network segmented into VLANs that separate VoIP, workstations and printers, employee wireless devices, surveillance cameras and recorders, and guest internet access. Its critical assets, all rated high criticality, are the Active Directory domain controller, the PACS medical-imaging server and the OsiriX DICOM web server that exposes imaging to clinicians and patients, and the Synology NAS that holds administrative, financial, and patient data. The sector's regulatory drivers are GDPR, the NIS2 Directive, and DORA, and its principal threats are ransomware and data breaches, vulnerabilities in Internet of Medical Things (IoMT) devices, and insider error compounded by limited staff awareness.

Configuration. For the pilot, the monitoring framework draws evidence from the components deployed in the AMC environment: Wazuh and Suricata for endpoint and network detection across the VLAN-segmented

© CYberSynchorny, 2025

infrastructure, the AMC Active Directory as the identity source (federated through Keycloak for the controls that depend on it), the Wazuh vulnerability detector for vulnerability evidence, MISP and Kitsune for threat-intelligence enrichment, CROSS-CORE and the Evidence Integrity System for secure data exchange and integrity proofs over the NAS and PACS data, and Cybergate & TrustSense for staff training. CYBERGOPLUS consumes these outputs and evaluates them against the controls most relevant to the healthcare assets: asset inventory across the VLANs (AM2), identity and access on patient-data systems (IAM1, IAM4), network segmentation and detection (NS4, NS7), vulnerability status on clinical and IoMT systems (VM1), backup and integrity of the NAS and PACS (DS3, DS), incident handling (EIM1, EIM2), staff awareness (TA2), and the governance controls (GOV1, GOV3, RM1).

Workflows exercised. Given the pilot's threat profile and its current awareness gap, the phishing workflow (Section 7.2) is the most prominent, exercising detection, enrichment, identity remediation on Active Directory accounts, and retraining. The vulnerability workflow (Section 7.3) applies to the clinical and IoMT systems, where legacy devices carry the greatest exposure. The identity workflow (Section 7.4) verifies access to the PACS, OsiriX, and NAS, where patient-data confidentiality depends on role-based access. The resilience and integrity workflow (Section 7.6) applies to the NAS bidirectional backup to OneDrive and to the integrity of the PACS imaging and DICOM data. The asset-reconciliation trace (Section 7.1.1) is significant here, because the environment spans many device classes across seven VLANs, including imaging machines and surveillance devices that cannot host agents.

Compliance posture. The pilot begins from an uneven posture. Several technical controls are already in place: VLAN isolation and firewall rules support the network controls; Active Directory enforces password complexity and periodic change, with no administrative rights for ordinary users; the NAS provides role-based access, bidirectional backup, and file versioning; and endpoints carry centrally managed antivirus with regular patching. The governance and awareness side is the starting weakness: AMC holds ISO 9001, ISO 14001, and ISO 45001 certifications but no information-security management certification, and at baseline it has no cybersecurity awareness program in place. The framework therefore measures the Identity Control, Asset Security, Threat Protection, Vulnerability Status, and Resilience areas from the technical evidence above, and the Governance and Strategy area from the policy and awareness evidence, with maturity computed as described in Section 5.3.3. No measured maturity figure is asserted here; the posture is expressed as the controls in scope and the evidence sources that populate them.

Operational considerations. Healthcare imposes specific constraints on monitoring. Medical imaging machines, the OsiriX macOS host, and IoMT devices often cannot run a monitoring agent, so their compliance evidence is derived from network-level observation through Suricata rather than from host telemetry, and asset reconciliation must treat them as managed but agentless. Patient data is special-category data under GDPR, so the evidence the framework collects from PACS, OsiriX, and the NAS must itself be minimized and protected, a consideration carried into Chapter 9. Availability is paramount in a clinical setting, and the use case explicitly has medical staff report any delay or disruption caused by security controls, so monitoring and any automated response must avoid interrupting care. Finally, the bidirectional NAS backup to OneDrive is both a resilience strength and a data-egress point that the secure-exchange and integrity controls must cover, and the guest Wi-Fi VLAN remains a deliberately isolated but ever-present part of the attack surface.

8.2 UC2: Public Administration (EGL)

Pilot profile and sector context. MAESTRO is an incident-management platform used by the Municipality of Egaleo (EGL) to handle Municipal Police incidents and coordinate emergency response across police, fire, emergency health, and other municipal authorities, and it is the public-administration use case. MAESTRO is not itself a security tool, but as critical operational infrastructure it is a high-value target: a successful attack could disrupt essential public services, paralyze real-time decision-making, or endanger citizens during an emergency. The regulatory drivers are data protection for the sensitive incident and coordination data and the immutability of records for legal and forensic purposes; the principal concerns are the operational dependency on the

© CYberSynchorny, 2025

platform, resilience to advanced threats, real-time monitoring and forensic readiness, and the awareness gap among operators and municipal staff.

Configuration. MAESTRO is a microservices platform built around a central MAESTRO Core, with a Consumers and Scheduler tier, a RabbitMQ message broker, MariaDB for structured records, logs, and event histories, MongoDB for unstructured feeds, message logs, and metadata, a Redis cache, REST and WebSocket APIs with InertiaJS, an Android app for field personnel, a push-notification system, and external and third-party integrations, all hosted in a high-availability private cloud. The monitoring framework integrates around this as a defensive layer that does not alter the core. The pilot configuration is based on the Wazuh, Suricata, Anomaly Detector, Kitsune, DCO, and CROSS-CORE components: Wazuh and Suricata, together with the Anomaly Detector, provide detection over the cloud-hosted services; the Wazuh vulnerability detector and the platform's regular vulnerability assessments provide vulnerability evidence; identity-related evidence is drawn from authentication-event monitoring, focusing on suspicious account activity and account-compromise scenarios rather than on the deployment of additional authentication controls; the DCO playbooks provide automated response; Kitsune provides external threat intelligence, including the detection of simulated information-leakage events; and CROSS-CORE provides the blockchain integrity proofs over MAESTRO's dispatch logs, incident responses, and communication records. CYBERGOPLUS evaluates these outputs against the controls most relevant to the platform: incident handling (EIM1, EIM2), secure communication (NS4) and intrusion detection (NS7), authentication monitoring and access control (IAM1), vulnerability status (VM1), the integrity and backup of incident records (DS, DS3), business continuity of the emergency service (BCR2), staff awareness (TA2), and the governance controls (GOV1, GOV3, RM1).

Workflows exercised. Following the latest UC2 pilot definition, the headline scenario combines account-compromise detection with incident-record integrity. The resilience and integrity workflow (Section 7.6) is the most prominent here: the pilot evaluates the detection of unauthorized modification or deletion of incident records, blockchain-based integrity verification through CROSS-CORE, and the detection of simulated information-leakage events through Kitsune, directly supporting the objective of tamper-proof public records, alongside the continuity of an emergency-critical service. The phishing workflow (Section 7.2) addresses the operator awareness gap and the human-error risk. Identity-related activity concentrates on authentication monitoring rather than on the deployment of authentication controls: it monitors authentication events, detects suspicious account behavior, and evaluates response capabilities following simulated account-compromise scenarios, drawing on the detection and incident-handling workflows. Vulnerability status is monitored continuously through the Wazuh vulnerability detector and the platform's regular vulnerability assessments. The detection path that opens the phishing workflow also runs continuously over the private-cloud services to satisfy the real-time monitoring need.

Compliance posture. MAESTRO begins from a relatively strong technical baseline on confidentiality and transport but a weaker position on monitoring and governance. Encryption is well covered: TLS protects data in transit, AES protects the MariaDB and MySQL stores, MongoDB Enterprise uses AES256-CBC and AES256-GCM, passwords are hashed with bcrypt, and the chat system runs over HTTPS and WSS, all of which supports the secure-communication control. Private-cloud isolation, rate limiting, and regular vulnerability testing further support the network and vulnerability areas. The gaps the use case sets out to close are real-time detection and forensic monitoring, which the framework's detection and integrity evidence addresses, and awareness, where no program yet exists for the target audience of first responders and municipal police employees. The framework measures the Threat Protection, Identity Control, Vulnerability Status, and Resilience areas from the technical evidence, and the Governance and Strategy area from the policy and awareness evidence, with maturity computed per Section 5.3.3 and no measured figure asserted.

Operational considerations. The defining constraint is that MAESTRO is life-safety-critical: monitoring and any automated response must not disrupt emergency coordination, which is why the CYberSynchrony modules are integrated as a non-intrusive defensive layer around the core rather than within it. The immutability of incident and dispatch records is a legal and forensic requirement, so the integrity workflow and the CROSS-CORE proofs are central rather than peripheral. Private-cloud isolation reduces the external attack surface but does not

© CYberSynchorny, 2025

remove the insider and human-error risk that the awareness gap leaves open. The platform's breadth, with external API endpoints to national services and identity providers, third-party integrations, and a mobile app for field staff, widens the surface that asset and identity monitoring must account for. As in the healthcare pilot, the sensitive personal and incident data the framework observes must itself be minimized and protected, a consideration carried into Chapter 9.

8.3 UC3: Academia (UPRC)

Pilot profile and sector context. UC3, led by the University of Piraeus Research Centre (UPRC), focuses on academic institutions, which are open, dynamic environments with a large and diverse user base, a need for seamless access to digital resources, and large volumes of sensitive data. The University of Piraeus environment supports roughly 29,000 active students and 500 instructors across nine undergraduate and more than 40 master's programs. Its critical assets, all rated high, are the Open eClass learning management system, the in-house Student Preregistration system, the Student Information System holding academic records and grades, and the Network Operations Center providing DNS, DHCP, and FreeRADIUS-based authentication. The principal threats are phishing, ransomware, and unauthorized access, and the interconnected nature of academic networks extends that risk to collaborative research and industry partnerships. Unlike the healthcare and public-administration sectors, academia is not a core NIS2 sector with an established adoption proxy, as noted in Annex II, so the pilot's baseline is set from the project's own assessment.

Configuration. The university operates a heterogeneous estate: Open eClass on Debian with Apache, PHP, and MariaDB; the Preregistration system on Ubuntu with WildFly, Spring Boot, and PostgreSQL; the three-tier Student Information System with SQL Server, Redis, and JasperServer; and NOC core services on ISC Bind, ISC KEA, and FreeRADIUS, fronted by Sophos XG and Cisco ASA with Firepower IPS over a Cisco and Huawei switching fabric, with a Windows Server 2019 Hyper-V failover cluster underneath. The monitoring framework draws evidence from the components deployed across this estate: Wazuh and Suricata for endpoint and network detection, the latter complementing the existing Firepower IPS; the university's AAA and single-sign-on as the identity source (FreeRADIUS and CAS, federated through Keycloak for the controls that depend on it); SNYK and the periodic vulnerability scans for vulnerability evidence, including web-application checks on the LMS and SIS; MISP and Kitsune for threat-intelligence enrichment; CROSS-CORE and the Evidence Integrity System for record integrity; and Cybergate for awareness training. CYBERGOPLUS evaluates these outputs against the controls most relevant to the academic estate: asset inventory across the network zones (AM2), identity and access on student and faculty systems (IAM1, IAM4), network detection and secure communication (NS4, NS7), vulnerability status on the public-facing web applications (VM1), incident handling (EIM1, EIM2), record integrity and backup (DS, DS3), awareness (TA2), and the governance controls (GOV1, GOV3, RM1).

Workflows exercised. The pilot's own test scenarios map directly onto the workflow chapter. The phishing workflow (Section 7.2) corresponds to the phishing scenario the research team defines and the CYBRITE module leader automates. The vulnerability workflow (Section 7.3) covers the SQL-injection scenario against the LMS and SIS web applications, where SNYK and the periodic web-application scans provide the evidence. The detection path that opens the phishing workflow also addresses the distributed-denial-of-service scenario against the NOC's DNS, where Suricata and Wazuh provide the telemetry. The identity workflow (Section 7.4) verifies the FreeRADIUS AAA and CAS single-sign-on access across the large user base, including MFA where supported. The asset-reconciliation trace (Section 7.1.1) is relevant across the multi-zone, multi-vendor network.

Compliance posture. Of the three pilots, the university begins from the most developed technical baseline. Next-generation firewalls and an intrusion-prevention system support the network controls; FreeRADIUS AAA with strong password policies and MFA where supported supports the identity controls; periodic vulnerability scanning with patch management supports the vulnerability area; and endpoint protection with regular OS patching is in place across a heterogeneous server estate. It is also the only pilot with an awareness program already in place, although a basic one: campaigns run for administrative staff and IT professionals, using internal emails and security checklists at a single proficiency level, with no audience-specific tailoring yet. The framework therefore measures the Threat Protection, Identity Control, Vulnerability Status, and Resilience areas from the

© CYberSynchorny, 2025

technical evidence, and the Governance and Strategy area from the policy and awareness evidence, with maturity computed per Section 5.3.3 and no measured figure asserted. The pilot's own KPIs, such as a targeted reduction in detected-and-resolved incidents and a faculty-compliance target after training, sit alongside the framework's indicators.

Operational considerations. The defining characteristic of the academic environment is openness at scale: tens of thousands of users, public-facing web applications, and a need for frictionless access that limits how aggressively access controls and automated responses can be applied. The estate is highly heterogeneous in operating systems, application servers, and databases, so monitoring coverage and asset reconciliation must span many platform types. The public-facing LMS, Preregistration, and SIS applications are the most exposed to the web-application attacks the pilot tests, while the NOC's DNS service is the focus of the availability scenario, making both detection latency and resilience important. Student and faculty records are personal data under GDPR, so, as in the other pilots, the evidence the framework collects must itself be minimized and protected, a consideration carried into Chapter 9. Finally, academia's position outside the core NIS2 sectors means the pilot also serves as a reference point for extending the framework to environments without an established sectoral compliance baseline.

8.4 Cross-Pilot Observations

Read together, the three pilot narratives show a single, sector-agnostic framework applied to three very different critical environments. The observations below draw out what the pilots share, where they diverge, and how the framework accommodates that divergence without changing its structure, which is the essence of the harmonization objective of Task 6.6.

A common starting posture. Despite their differences, all three pilots begin from an uneven baseline in which technical controls are partially present while governance and awareness lag. The awareness dimension is the clearest shared weakness: the healthcare and public-administration pilots have no cybersecurity awareness program at baseline, and the academic pilot has only basic campaigns for a narrow audience. None of the three holds an information-security management certification at the outset. This is consistent with the sectoral baseline evidence in Annex II, where documented governance lags behind technical hygiene.

A maturity gradient. The pilots sit at different points on a technical-maturity gradient. The university presents the most developed baseline, with next-generation firewalls, intrusion prevention, centralized AAA, multi-factor authentication where supported, and routine vulnerability scanning. The public-administration platform is strong on confidentiality and isolation, with private-cloud deployment and pervasive encryption, but weaker on detection and forensic monitoring. The healthcare center has sound basic hygiene, with VLAN segmentation, restricted Active Directory access, and managed endpoints, but the least formal program. The framework expresses this gradient through the same maturity model rather than through separate scales.

Sectoral emphases. Each pilot stresses a different cluster of controls. Healthcare centers on patient-data confidentiality, IoMT exposure, and clinical availability. Public administration centers on the immutability of incident and dispatch records for legal and forensic purposes and on the continuity of a life-safety service. Academia centers on openness at scale, public-facing web applications, and the availability of core network services such as DNS. The framework accommodates these emphases by varying which controls are weighted and which evidence sources are in scope, not by altering the control catalogue.

Abstraction over tools. A single control is populated from different evidence sources across the pilots. Identity and access control (IAM1), for example, is evidenced from Active Directory in the healthcare pilot, from the MAESTRO application's authentication layer in the public-administration pilot, and from FreeRADIUS and CAS single-sign-on in the academic pilot, yet it maps to the same control and indicator. This abstraction over tools is what allows heterogeneous environments to be compared on a common basis.

Harmonization. Taken together, these observations demonstrate the harmonized monitoring approach the project set out to validate. The same control framework, the same KCIs, the same maturity model, and the same

© CYberSynchorny, 2025

workflows apply across healthcare, public administration, and academia, with sectoral difference handled through configuration. This common basis is also what makes the cross-sector adoption measure of KPI 9.2 meaningful, since the three pilots are assessed against one consistent set of controls and indicators rather than three incompatible ones.

© CYberSynchorny, 2025

9 Ethics, Privacy and Data Protection

The compliance monitoring framework exists to hold an organization's security controls to account, but in doing so it observes and processes data that can itself be personal, including authentication and login records, role and access assignments, and training outcomes. A monitoring tool that enforces data-protection controls must meet the same standard it measures. This chapter turns that scrutiny inward, examining the ethical, privacy, and data-protection implications of CYBERGOPLUS itself: the categories of personal data it processes, the lawful basis and data-minimization principles that govern that processing, the technical and organizational measures that protect it, and the privacy-by-design choices built into the framework and its evidence sources. The analysis is framed by the GDPR and by the data-protection obligations common to the three pilot sectors, and it applies to the framework the same level of examination that the project applies to the systems it monitors.

9.1 Personal Data Processed and Lawful Basis

The framework computes compliance from evidence, and a defined subset of that evidence is personal data. A precise inventory of what is processed, and why, is the starting point for the data-protection analysis. One point frames the whole chapter: CYBERGOPLUS processes security and compliance metadata, not the underlying sensitive content of the systems it monitors. It records that an authentication occurred and whether it used multi-factor authentication, not the contents of a patient record; that a training module was completed, not the substance of a consultation; that an asset was observed on the network, not the citizen data an incident concerned. The personal data it does process falls into the categories set out in the table below.

Table 12: Personal Data Categories Processed by CYBERGOPLUS

Data Category	Primary Source	Purpose (Control)	Lawful Basis
Authentication and login records	Keycloak login events and token claims	Verify strong authentication (IAM4)	Legitimate interest; legal obligation
Identity and role assignments	Keycloak roles; Role Assignment Matrix	Verify role-based access and accountability (IAM1, GOV1)	Legitimate interest; legal obligation
Training records and outcomes	Cybergate mission and simulation results	Verify awareness training (TA2)	Legitimate interest; legal obligation
Incident-related identifiers	Phishing and incident workflow records	Verify incident handling and remediation (EIM2)	Legitimate interest; legal obligation
Network and asset identifiers, including IP addresses	Wazuh and Suricata telemetry	Verify asset inventory and detection (AM2, NS7)	Legitimate interest
Findings and remediation records	CYBERGOPLUS finding workflow	Track and close non-compliance (REQ-FUNC-07, REQ-FUNC-08)	Legitimate interest; legal obligation

The processing rests on two lawful bases under Article 6 of the GDPR. The primary basis is the legitimate interest of the controller in the security of its network and information systems, which Recital 49 of the GDPR expressly

© CYberSynchorny, 2025

recognizes as a legitimate interest. The second basis is legal obligation, where the NIS2 Directive and the sector-specific obligations applicable to each pilot require the organization to monitor, evidence, and report on these controls. Where the data concerns employees, such as login and training records, the processing is limited to what is necessary to verify the controls and is not used for individual performance evaluation, a boundary that the proportionality discussion later in this chapter revisits.

The framing above addresses employee data, but in the academic use case the majority of monitored identities are students rather than staff. The lawful basis is the same in kind: the controller's legitimate interest in the security of its network and information systems (Article 6(1)(f) and Recital 49 of the GDPR) covers the authentication and access evidence generated by student accounts, and the legal-obligation basis applies where regulatory duties require it. The distinction is one of purpose rather than basis: for students, the evidence verifies the security of access to learning and information systems, and it is likewise not used to evaluate or profile individuals, consistent with the proportionality limits set out later in this chapter.

Two further points complete the inventory. The trustSense team assessment, although it supports the risk controls, is excluded from this inventory because it retains no personal data on its users or on assessment outcomes, a privacy-by-design property discussed later in this chapter. And special-category and other protected data, in particular the patient data in the healthcare pilot, the citizen incident data in the public-administration pilot, and the student and faculty academic and personal records in the academic pilot, is outside the scope of the framework's own processing by design, since the framework consumes the metadata of the controls protecting that data rather than the data itself.

9.2 Data Minimization and Protective Measures

The data inventoried in Section 9.1 is governed by data-minimization principles and protected by a set of technical and organizational measures. The guiding principle is that the framework applies to itself the controls it measures in others: it processes the least data necessary, protects it as it would require any monitored system to protect equivalent data, and verifies the integrity of the evidence it relies on.

Data minimization and purpose limitation. The framework collects only the data each KCI requires, which is typically a flag, a timestamp, and an identifier rather than the content behind them. Identities are represented by their Keycloak subject identifier wherever the monitoring logic permits, so that evidence and findings can be processed without resolving to a named individual. The data is used solely to verify the controls and to track remediation, and not for individual performance evaluation.

Retention. Evidence records and findings are retained only as long as necessary to verify the relevant control and to satisfy the organization's audit and reporting obligations. Once a finding is closed and the applicable retention period has passed, the underlying personal-data records are purged or reduced to the aggregate KCI and metric history, which carries no personal data and is what feeds the longer-term trend and maturity views.

Access control and least privilege. The framework does not hold its own credentials; it federates authentication to Keycloak, and the presentation API scopes every response to the roles asserted in the caller's token, so that a user sees only the compliance data their role permits. The ingestion service account is restricted to the minimal administrative roles its connectors require, such as the view-users and view-events roles on the identity realm, and the API gateway records access for auditability.

Encryption and secure exchange. All communication is encrypted in transit using TLS 1.2 or higher, and stored evidence is encrypted at rest. Intelligence and evidence exchanged through the Secure Data Exchange Platform inherit its v2 cryptographic protections, including AES-GCM encryption, RSA-OAEP key exchange, and SHA-256 integrity hashing, and are shared under controlled access with TLP taxonomy tags.

Integrity and tamper-evidence. The evidence the framework relies on is itself made verifiable: critical logs and records are anchored to the Evidence Integrity System through the CROSS-CORE interface, so that any alteration of an evidence record is detectable against its stored blockchain hash. This protects the integrity of the compliance record as strongly as the framework requires of the systems it monitors.

© CYberSynchorny, 2025

Taken together, these measures mean the monitoring tool meets the data-protection and integrity controls it enforces. The framework's own processing is minimized, access-controlled, encrypted, time-bound, and tamper-evident, which is the standard it holds the monitored systems to.

9.3 Privacy by Design and by Default

Several of the choices described in the preceding sections are not incidental; they are deliberate applications of the data-protection-by-design-and-by-default principle of Article 25 of the GDPR. This section gathers them, beginning with the clearest example.

Team-level assessment with no data retention (trustSense). The trustSense risk assessment is built so that it retains no personal data on its users or on assessment outcomes. It produces a team-level trustworthiness score and a set of recommended mitigations without storing who said what, so the human-and-organizational risk control is evidenced at the level of the team rather than the individual. This is privacy by design in its strongest form: the data that could identify an individual is not collected in the first place.

Metadata rather than content. As established in Section 9.1, the framework consumes the metadata of the controls it verifies, not the sensitive content those controls protect. By design, patient records, citizen incident data, student & faculty records, and the substance of communications never enter the framework. The scope of processing is constrained at the architectural level, not by a downstream policy that could later be relaxed.

Aggregation by default. The framework's default output is aggregate. KCIs roll up into rates, KPIs, and maturity levels, and the board-level dashboard presents this aggregate, non-personal view. Personal-level detail, such as the identity attached to a specific finding, is generated only where a control deviation requires remediation and is exposed only to the roles authorized to act on it. The default view is therefore the least revealing one.

Pseudonymization and federated identity by default. Identities are represented by their Keycloak subject identifier rather than by name wherever the logic permits, and the framework holds no credentials of its own, federating all authentication to Keycloak. The default state is that the framework processes pseudonymous references and never becomes a second store of identity secrets.

These choices implement data protection by design and by default rather than as an overlay. The framework's least-revealing configuration is its normal configuration, and the most sensitive categories of data are excluded at the level of the architecture rather than controlled after collection.

9.4 Ethics and Proportionality

Beyond its strict data-protection obligations, the framework raises ethical considerations that follow from its position: a tool that continuously observes the controls, and therefore in part the behavior, of an organization and its staff. This section sets out how proportionality, transparency, and human oversight are maintained.

Proportionality of monitoring. The framework observes authentication, access, and training data, which are also records of how individual staff behave. Proportionality requires that this observation stays tied to its security and compliance purpose and does not become general workplace surveillance. The boundary stated in Section 9.1 is firm: this data verifies controls and is not used to evaluate individual performance. The team-level design of the trustSense assessment reinforces this, evidencing the human-risk control without ranking individuals.

Transparency. Staff whose data is processed are entitled to know that compliance monitoring takes place and what it covers. The categories in Section 9.1 are intended to be reflected in the organization's privacy notices, so that monitoring is disclosed rather than covert, consistent with the transparency obligations of the GDPR.

Purpose limitation and function creep. The data is collected to verify controls, and the framework is designed to resist its reuse for other ends. Aggregation by default, role-scoped access, and time-bound retention all narrow the opportunity for the compliance record to be repurposed as, for example, a productivity or attendance record.

© CYberSynchorny, 2025

Human oversight of automated findings. The framework automatically generates findings when a control deviates, but these are findings about controls and assets, not automated decisions that produce legal or similarly significant effects on individuals. Where a finding concerns a person, such as an account without multi-factor authentication, the response is a security remediation rather than an adverse judgment, and the remediation workflow keeps a human in the loop for assignment and closure.

Proportionate response. Where the workflows trigger an automated action, such as forcing a password reset or isolating a host, the action must be proportionate to the threat and must not disrupt an essential service. As the pilot narratives noted, this constraint is strongest in the life-safety and clinical settings, where monitoring and response are deliberately bounded so as not to interrupt care or emergency coordination.

The framework's ethical posture follows from the same principle as its data protection: it is built to do the least necessary to achieve a legitimate security purpose, transparently, proportionately, and with human judgment retained where it matters. These commitments, and the limits that remain, are revisited in the limitations reflection of Chapter 10.

© CYberSynchorny, 2025

10 Conclusions

This deliverable has developed the compliance monitoring framework from a defined capability into an operational, evidence-driven system. Building on the foundations established in D6.4, it has specified the end-to-end workflows that carry evidence from the integrated components through the CYBERGOPLUS engine to compliance findings and remediation, set out a measurement methodology that grounds each indicator in an explicit baseline, target, and evidence source, deepened the integration interfaces to the level of concrete endpoints and data models, applied the framework to the three sectoral pilots as integrated narratives, and held the framework to its own data-protection and ethical standards. This chapter draws those contributions together, states honestly what the framework does and does not yet demonstrate, reflects on the principal risks to the approach, and sets out the path toward full validation.

This progression from blueprint to operational specification is the central advance of the version. The framework is no longer described only in terms of its structure and architecture; it is specified in terms of how it operates. The end-to-end workflows of Chapter 7, the concrete integration interfaces and data models of Chapter 6 and Annex I, and the per-pilot operational narratives of Chapter 8 are the elements that carry the framework across that line, from a design that can be built to a design specified in operational detail.

10.1 Summary of Contributions

Across the preceding chapters, this deliverable advanced the compliance monitoring framework along several dimensions, each building on the structure and content of D6.4 rather than replacing it. The principal contributions are summarized below.

- **An expanded evidence base.** Chapter 5 extended the framework's evidence sources with MISP for threat-intelligence enrichment and secure sharing, the upgraded Kitsune for external intelligence and structured incident management, and trustSense for team-level risk assessment, and it consolidated the operational indicators into a single KCI catalogue and comprehensive indicator map.
- **A measurement methodology.** The methodology of Section 5.3.4 grounds each reported indicator in an explicit baseline, target population, measurement window, and evidence source, and it established the triangulated baseline for the industry-specific standard-adoption indicator (KPI 9.2), with the supporting evidence consolidated in Annex II.
- **Deeper integration specifications.** Chapter 6 specified the integration interfaces to the level of concrete endpoints, authentication models, and data formats, adding the MISP and Kitsune interfaces alongside the existing Wazuh, CROSS-CORE, and Keycloak ones, and surfacing the technical evolution of the secure-exchange and evidence-integrity components since D6.4.
- **End-to-end workflows.** Chapter 7 set out the cross-module workflows that carry evidence from generation through validation to finding and remediation, traced one workflow to the level of the API calls, and confirmed through a coverage matrix that every control in the framework is exercised.
- **Per-pilot integrated narratives.** Chapter 8 applied the framework to the three sectoral pilots as coherent end-to-end narratives and drew them together to demonstrate the harmonized, sector-agnostic design.
- **A data-protection and ethics analysis.** Chapter 9 held the framework to its own standard, inventorying the personal data it processes, the protective measures and privacy-by-design choices that govern it, and the ethical and proportionality boundaries it observes.

Together, these contributions move the framework from a described capability to an operational and auditable system, while keeping its conceptual scope as a meta-tool that consumes the documented outputs of the integrated components.

© CYberSynchorny, 2025

10.2 Limitations

The framework's scope and stage of development impose limitations that should be read alongside its contributions. Stating them plainly is part of presenting an honest account of what the deliverable establishes.

Methodology-forward, not yet longitudinally validated. In line with the project timeline, this deliverable defines and configures the monitoring framework rather than reporting its operation over time. Full pilot deployment and longitudinal validation are not demonstrated by Month 36, and the indicators are presented as a measurement approach rather than as measured outcomes. The validation framework is defined; the longitudinal evidence will accrue beyond submission.

Dependence on the source components. CYBERGOPLUS is a meta-tool that consumes the documented outputs of the integrated components, so its assessment is only as complete and accurate as those sources. A control that a source tool fails to detect, or a feed that is stale or unavailable, is a gap the framework can surface as missing evidence but cannot independently correct. The framework verifies that the controls report, not that the underlying tools are themselves infallible.

Compliance is not the same as security. The framework verifies that defined controls are present and operating, expressed against the regulatory framework. A high compliance posture indicates that the controls in scope are working, not that the organization is immune to every threat, including those outside the control catalogue. The output should be read as evidence of control operation, not as a guarantee of security.

Partial coverage of agentless and legacy assets. As the pilot narratives noted, some assets cannot host a monitoring agent, including medical imaging machines, certain IoMT devices, and life-safety platforms that must not be disturbed. These are observed at the network level rather than from host telemetry, which yields less granular evidence for the controls that depend on host data.

Proxy-based KPI baseline. The baseline for the industry-specific standard-adoption indicator (KPI 9.2) is a triangulated proxy drawn from recognized external indicators, not a direct compliance count. It is a defensible starting point, and it will be superseded by the project's own survey of the use-case participants, but until then it carries the uncertainty inherent in any proxy.

Threshold and configuration sensitivity. The indicators depend on thresholds and on which evidence sources are in scope for each environment. Thresholds set too loosely risk false assurance, and too tightly risk noise. Calibration is a configuration responsibility that the framework supports but does not remove.

These limitations do not undermine the framework; they delimit it. Each points to where care is needed in deployment and where the risks discussed in the next section concentrate.

10.3 Risk Reflection

Distinct from the inherent limitations above, several risks could affect the framework as it moves toward deployment. Each is stated with the mitigation that contains it, and each is consistent with the project risk register, which carries the active WP6 entries and the retrospective WP2 entry.

External tooling integration drift. The framework depends on the interfaces of components it does not control, and external or rapidly evolving tools, such as Kitsune and the analytics tooling around WEKA, may change their APIs in ways that break a connector. This is mitigated by the thin-connector design, in which each source is reached through a standard client library behind the API gateway, so that an interface change is isolated to one connector and surfaced quickly through connector-health monitoring rather than propagating into the engine.

Evidence-source degradation. If a single source degrades, goes offline, or returns stale data, the controls that depend on it lose their evidence. The framework treats this fail-safe: absent or stale evidence is recorded as a missing-evidence finding rather than assumed to pass, so a degraded source lowers the posture and raises a finding instead of silently inflating it. Where an assessment area draws on more than one source, the loss of one does not blind the whole area.

© CYberSynchorny, 2025

Timeline pressure. D6.8 and the components it builds on are on a compressed schedule toward Month 36, and integration specification can outpace deployment. The methodology-forward framing is itself the mitigation: by defining and configuring the validation framework first, the project ensures that deployment and longitudinal measurement can proceed on a defined basis after submission rather than being rushed before it.

Threshold mis-calibration. As noted among the limitations, indicator thresholds that are set too loosely or too tightly distort the posture. The mitigation is to make thresholds explicit and configurable, to start them conservatively, and to revisit them on the same review cycle as the governance artifacts, so that calibration improves with operational experience.

Misinterpretation of scope. There is a risk that the framework is read as an audit of pilot deployments or as a guarantee of security. The mitigation is the consistent framing maintained throughout this deliverable: CYBERGOPLUS is a conceptual meta-tool that evidences the operation of defined controls, and a compliance posture is evidence of control operation, not a security guarantee.

None of these risks is novel to this framework; each is a familiar risk of integration-heavy monitoring systems, and each is contained by a design choice already described in this deliverable rather than deferred to deployment.

10.4 Outlook and Path to Validation

The framework is defined and configured; what remains is to operate it and to accumulate the evidence that turns its measurement approach into measured results. The path to full validation is clear and proceeds beyond the Month 36 submission.

From configuration to operation. The immediate next step is to run the configured workflows in the pilot environments over time, executing the monitoring loop continuously so that the indicators presented here as a measurement approach become measured values. This converts the designed KCIs, metrics, and maturity levels into an operational record for each pilot.

Populating the standard-adoption baseline. The proxy baseline for KPI 9.2 will be replaced by the project's own survey of the use-case participants, which provides a direct measure of standard adoption to assess against the 65% target. That result is reported in D8.3, at which point the indicator rests on a primary measurement rather than a triangulated proxy.

Longitudinal validation. As the KCI, metric, and maturity history accumulates, it will validate that the indicators behave as designed and that the workflows operate end-to-end in live conditions, including the cross-module sequences that are most difficult to exercise in a single environment. This longitudinal evidence is the principal item that the project timeline defers beyond submission.

Calibration with experience. Operating the framework provides the data needed to calibrate thresholds, refining them on the governance review cycle so that the indicators converge on a balance between false assurance and noise. This directly addresses the calibration risk identified above.

Broader applicability. The harmonized, sector-agnostic design positions the framework for use beyond the three pilots. The academic pilot, sitting outside the core NIS2 sectors, serves as a reference point for extending the approach to environments that lack an established sectoral compliance baseline, supporting the project's wider objective of cross-sector applicability.

In sum, this deliverable leaves the framework operational in design and ready for deployment. The contributions are in place, the limitations and risks are understood and mitigated, and the remaining work is the longitudinal validation that the project timeline appropriately places after submission.

© CYberSynchorny, 2025

11 References

[1] Author, Title, Date, URL or page address European Parliament and Council, **Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union (NIS 2 Directive)**, 14 December 2022. URL: http://data.europa.eu/eli/dir/2022/2555/oj

[2] European Parliament and Council, Regulation (EU) 2016/679 on the protection of natural persons with regard to the processing of personal data (GDPR), 27 April 2016. URL: http://data.europa.eu/eli/reg/2016/679/oj

[3] International Organization for Standardization, ISO/IEC 27001:2022 Information security, cybersecurity and privacy protection — Information security management systems — Requirements, October 2022. URL: https://www.iso.org/standard/27001

[4] Digital Security Authority (DSA), **The Security of Networks and Information Systems (Security Measures of Operators of Essential Services and Operators of Critical Information Infrastructure) Decision, 2020 (K.D.P. 389/2020)**, August 2020. URL: https://dsa.cy/en/legislation/decisions

[5] National Institute of Standards and Technology (NIST), **Security and Privacy Controls for Information Systems and Organizations (SP 800-53 Rev. 5)**, September 2020. URL: https://csrc.nist.gov/pubs/sp/800/53/r5/final

[6] European Union Agency for Cybersecurity (ENISA), "Cybersecurity Standards and National Authorities". URL: https://www.enisa.europa.eu/audience/national-eu-authorities

[7] CYberSynchrony Deliverable "D2.2 Cybersecurity landscape and threat analysis end-user requirement analysis and use case definition and CYberSynchrony Framework", Submitted M16

[8] CYberSynchrony Deliverable "D3.2 Risk Scenario Analysis Toolkit", Submitted M16

[9] CYberSynchrony Deliverable "D4.1 Integrated Cybersecurity Monitoring Suite", Submitted M16

[10] CYberSynchrony Deliverable "D4.2 Automated Response Toolkit", Submitted M12

[11] CYberSynchrony Deliverable "D5.2 Evidence Integrity System", Submitted M13

[12] CYberSynchrony Deliverable "D6.1 Cyber Awareness Platform", Submitted M19

[13] CYberSynchrony Deliverable "D6.2 Gamified Training Modules", Submitted M19

[14] CYberSynchrony Deliverable "D6.3 Cybersecurity Policy Toolkit", Submitted M19

[15] CYberSynchrony Deliverable "D8.1 CYberSynchrony Integration Blueprint Use Case Portfolio", Submitted M16

[16] CYberSynchrony Deliverable "D5.3 Secure Data Exchange Platform v2", Submitted M25

[17] CYberSynchrony Deliverable "D5.4 Evidence Integrity System v2", Submitted M25

[18] CYberSynchrony Deliverable "D4.3 Integrated Cybersecurity Monitoring Suite v2", Submitted M25

[19] CYberSynchrony Deliverable "D3.3 Enhanced Threat Intelligence System v2", Submitted M25

[20] CYberSynchrony Deliverable "D6.5 Cyber Awareness Platform v2", To be Submitted M27

[21] CYberSynchrony Deliverable "D6.6 Gamified Training Modules v2", To be Submitted M27

[22] CYberSynchrony Deliverable "D6.7 Cybersecurity Policy Toolkit v2", To be Submitted M27

© CYberSynchorny, 2025

Annex I: Technical Interface Specifications

This annex provides the detailed technical specifications for the interfaces summarized in Chapter 6. For each evidence source, and for the presentation layer, it sets out the protocol, the authentication method, the specific endpoints, and the request and response data models that the CYBERGOPLUS connectors rely on. It is the reference companion to the integration design in Section 6.3.

I.1 Conventions

The specifications in this annex follow a common set of conventions.

- **Transport and format**: All interfaces use HTTPS and exchange JSON, except the CROSS-CORE Evidence Integrity interface, which uses RESTAPI over HTTPS as noted in its section. Path parameters are shown in braces, for example {realm} or {agent_id}.
- **Authentication patterns**: Four patterns recur and are described once here, then referenced per interface. The Wazuh Manager issues a JWT from a login endpoint that is then sent as a Bearer credential. MISP uses a static automation API key sent in the Authorization header. Kitsune uses a Bearer JWT. Keycloak and the CYBERGOPLUS presentation API use OpenID Connect, validating a Keycloak-issued JWT against the realm JWKS.
- **Schema presentation**: Each data model is given either as a field listing, with the field name, type, and meaning, or as a representative JSON object, whichever is clearer for the interface. Representative objects show structure and are not exhaustive of every optional field.
- **Versioning caveat**: The specifications reflect the documented interfaces of the integrated components. Exact field availability and endpoint paths can vary with component version, so each connector validates the response shape it receives rather than assuming it.

The annex is organized by interface:

- **I.2**: Wazuh Manager and Indexer.
- **I.3**: Suricata EVE records.
- **I.4**: CROSS-CORE Evidence Integrity service.
- **I.5**: MISP REST API.
- **I.6**: Kitsune API.
- **I.7**: Keycloak identity services.
- **I.8**: CYBERGOPLUS presentation API.

I.2 Wazuh Manager and Indexer

Overview. The framework uses two Wazuh interfaces. The Wazuh Manager API (default port 55000) provides asset, vulnerability, configuration, and security-assessment data. The Wazuh Indexer (an OpenSearch service, default port 9200) provides the detection alerts, including the Suricata events forwarded through the agent. The Manager API is used for the asset, vulnerability, and configuration controls; the Indexer is queried for the detection controls.

Authentication. A client authenticates to the Manager with HTTP Basic credentials at POST /security/user/authenticate and receives a short-lived JWT in the response field data.token. The token is then sent as an Authorization: Bearer credential on all subsequent Manager calls. Indexer queries authenticate separately against the OpenSearch security layer.

© CYberSynchorny, 2025

Manager API response envelope. Manager responses share a common envelope: a data object containing affected_items (an array of result objects), total_affected_items, and total_failed_items, alongside message and error fields. The endpoints below are described by the fields within each affected item.

- **GET /agents** (asset inventory, AM2): each item carries id, name, ip, status (for example active or disconnected), os (with name, version, and platform), version, and lastKeepAlive.
- **GET /vulnerability/{agent_id}** (vulnerability status, VM1): each item carries cve, name (the affected package), version, architecture, severity (Low, Medium, High, or Critical), status, detection_time, and the cvss score.
- **GET /syscollector/{agent_id}/netaddr** (network addresses for asset reconciliation, AM2): each item carries iface, proto, address, netmask, and broadcast.
- **GET /sca/{agent_id}** (security configuration assessment, GOV3): each item carries policy_id, name, description, pass, fail, invalid, total_checks, and score.

Indexer detection alerts. Detection evidence is read from the wazuh-alerts-* indices with an OpenSearch query at POST /wazuh-alerts-*/_search, using a query that filters on the measurement window (a range on timestamp) and, where needed, on rule.level or group. The matching documents are returned under hits.hits, each with the _source structure shown below, which is the integrated Wazuh and Suricata record.

```
JSON
{
  "_index": "wazuh-alerts-4.x-2026.05.19",
  "_id": "<document id>",
  "_source": {
    "input": { "type": "log" },
    "agent": { "id": "002", "name": "cybersync4_100.64.12.14", "ip": "100.64.12.14" },
    "manager": { "name": "wazuh-manager-..." },
    "data": {
      "description": "Suricata: Alert - ET INFO DYNAMIC_DNS Query to nip.io Domain",
      "groups": ["ids", "suricata"],
      "id": "86601",
      "dns": { "query": [ { "rrname": "wazuh-dashboard...", "rrtype": "AAAA", "type": "query" } ] }
    },
    "location": "/var/log/suricata/eve.json",
    "decoder": { "name": "json" },
    "id": "1779188259.59417791",
    "timestamp": "2026-05-19T10:57:39.752+0000"
  }
}
```

© CYberSynchorny, 2025

For Suricata-sourced events, as above, the signature identifier and classification appear under data (data.id, data.groups, and the protocol fields such as data.dns), with location pointing at the Suricata eve.json. For Wazuh-native endpoint alerts, the _source additionally carries a rule object with rule.id, rule.level, rule.description, and rule.groups. The detection KCIs (EIM1 and NS7) read the rule level, the groups, the agent identity, and the timestamp from these documents.

I.3 Suricata EVE Records

Overview. Suricata writes events to /var/log/suricata/eve.json, one JSON object per line, each tagged with an event_type. The Wazuh agent forwards these into the wazuh-alerts indices described in Section I.2, where the EVE fields appear under data. Read directly from the file, the events have the structure below. The framework uses three event types: alert for detection, flow for asset reconciliation, and tls for secure-protocol verification.

Common envelope. Every EVE record carries a shared set of fields: timestamp, flow_id, in_iface, event_type, and the connection tuple src_ip, src_port, dest_ip, dest_port, and proto. The event-type-specific object follows.

alert (intrusion detection, NS7, and incident detection, EIM1). The alert object identifies the signature that matched and its classification. Indicator-based detections, where Suricata matches an indicator shared from MISP as described in Section I.5, appear as alert events whose signature corresponds to the shared indicator.

```
JSON
{
 "timestamp": "2026-05-19T10:57:38.473+0000",
 "flow_id": 1879188259594177,
 "in_iface": "eth0",
 "event_type": "alert",
 "src_ip": "100.64.12.14", "src_port": 51000,
 "dest_ip": "100.64.12.11", "dest_port": 53, "proto": "UDP",
 "alert": {
  "action": "allowed",
  "gid": 1,
  "signature_id": 2018452,
  "rev": 5,
  "signature": "ET INFO DYNAMIC_DNS Query to nip.io Domain",
  "category": "Potentially Bad Traffic",
  "severity": 2
 }
}
```

flow (asset reconciliation, AM2). The flow object summarizes a completed network flow. The src_ip and dest_ip from the envelope are the active network addresses that the asset-reconciliation trace of Section 7.1.1 compares against the Wazuh inventory.

The flow object carries pkts_toserver, pkts_toclient, bytes_toserver, bytes_toclient, start, end, and state.

© CYberSynchorny, 2025

tls (secure communication, NS4). The tls object records the negotiated parameters of a TLS session. The NS4 KCI reads tls.version and passes when the observed sessions use TLS 1.2 or higher.

```
JSON
{
 "timestamp": "2026-05-19T11:02:14.108+0000",
 "event_type": "tls",
 "src_ip": "100.64.12.20", "src_port": 49612,
 "dest_ip": "203.0.113.10", "dest_port": 443, "proto": "TCP",
 "tls": {
  "version": "TLS 1.3",
  "sni": "example.org",
  "subject": "CN=example.org",
  "issuerdn": "C=US, O=Let's Encrypt, CN=R3",
  "fingerprint": "1a:2b:..."
 }
}
```

Context event types. Two further event types provide context to an alert without being primary control evidence: dns records carry type, rrname, rrtype, and rcode; http records carry hostname, url, http_method, status, and length. They are retained for incident analysis (EIM2) but are not scored directly.

I.4 CROSS-CORE Evidence Integrity Service

Overview. The CROSS-CORE Evidence Integrity System (D5.2, advanced to v2 in D5.4) establishes the integrity of logs, backup archives, and shared intelligence by anchoring their cryptographic digests to a distributed ledger. CYBERGOPLUS uses it for one function: to verify that an evidence artifact has not been altered since it was anchored. This supports the integrity and backup controls (DS, DS3) exercised in the resilience workflow of Section 7.6.

Mechanism. As specified in D5.4, the system computes a Secure Hash Algorithm digest of an artifact and anchors that digest to a distributed ledger through a blockchain anchoring process. Verification recomputes the digest of the current artifact and compares it against the anchored record; a match confirms integrity, a mismatch indicates alteration. The Evidence Integrity System exposes RESTful APIs for exchange and uses MQTT and Apache Kafka for real-time distribution, and it integrates with MISP, the ECN, and the CSIRT Network for cross-organization sharing.

Verification interface. Consistent with the integration design in Section 6.3.1, CYBERGOPLUS invokes the integrity-verification function over HTTPs/REST, calling the CROSS-CORE API gateway's POST `/api/crosscore/verify/result` endpoint with a Keycloak-issued OIDC JWT as described in Section 6.3.2. The call re-hashes the evidence artifact and checks it against the anchored `resultHash`; the service returns whether the digest matches, together with the ledger transaction reference.

© CYberSynchorny, 2025

The request carries:

```
JSON
{
  "artifact_id": "log:wazuh-alerts-4.x-2026.05.19:FuniP54B",
  "artifact_hash":
"sha256:9f86d081884c7d659a2feaa0c55ad015a3bf4f1b2b0b822cd15d6c15b0f00a08",
  "artifact_type": "log"
}
```

The response carries:

```
JSON
{
  "artifact_id": "log:wazuh-alerts-4.x-2026.05.19:FuniP54B",
  "anchored_hash":
"sha256:9f86d081884c7d659a2feaa0c55ad015a3bf4f1b2b0b822cd15d6c15b0f00a08",
  "match": true,
  "tx_ref": "0x7b3f...e1",
  "anchored_at": "2026-05-19T10:58:02Z"
}
```

The artifact_type is one of log, backup, or shared_file, covering the three integrity uses: log integrity for detection evidence, backup-archive integrity for DS3, and shared-file integrity for the secure-exchange path. The integrity KCIs of Section 7.6 read the match flag and anchored_at: a false match raises a high-severity finding, and a stale or missing anchored record is treated as missing evidence per the fail-safe rule of Section 10.3.

I.5 MISP REST API

Overview. MISP (D5.3) provides threat-intelligence events and attributes for detection enrichment and records the secure-sharing evidence. The framework reads indicators for correlation against telemetry (EIM1), writes new indicators discovered during incident response (the phishing workflow of Section 7.2), and reads the sharing metadata, the TLP tags and distribution levels, for the secure-sharing control (DS).

Authentication. The MISP automation API key is sent in the Authorization header, with Accept and Content-Type both set to application/json. The key scopes the caller to the events and attributes its organization is permitted to see.

Endpoints.

- **POST /events/restSearch**: searches events using a JSON filter body and returns matching events under a response array.
- **POST /attributes/restSearch**: searches individual indicators and returns them under a response object containing an Attribute array, used to pull specific IPs, domains, URLs, and hashes for correlation.

© CYberSynchorny, 2025

- **GET /feeds**: lists the configured intelligence feeds with their enabled state, used to confirm feed currency.
- **POST /events**: creates an event, used to record new indicators identified during incident response.

The same search endpoints can return the results in different formats through the returnFormat parameter, including json, stix2, and the suricata and snort rule formats, so that intelligence can be pulled directly as detection rules.

Request. A representative detection-enrichment query, pulling published, detection-flagged indicators from the last seven days:

```
JSON
POST /events/restSearch
{
  "returnFormat": "json",
  "last": "7d",
  "published": true,
  "to_ids": true,
  "type": ["ip-dst", "domain", "url", "sha256"],
  "limit": 100
}
```

© CYberSynchorny, 2025

Response. Events are returned as objects under response, each wrapping an Event with its metadata, TLP tag, and attributes:

```
JSON
{
 "response": [
  {
   "Event": {
    "id": "1423",
    "date": "2026-05-18",
    "threat_level_id": "2",
    "analysis": "2",
    "info": "Phishing campaign targeting healthcare",
    "published": true,
    "uuid": "5f9c1e2a-...",
    "distribution": "3",
    "Tag": [ { "name": "tlp:amber" } ],
    "Attribute": [
     {
      "id": "98765",
      "type": "url",
      "category": "Network activity",
      "value": "http://malicious.example/login",
      "to_ids": true,
      "timestamp": "1747560000",
      "Tag": [ { "name": "tlp:amber" } ]
     }
    ]
   }
  }
 ]
}
```

Field semantics and control mapping. The Attribute value, type, and to_ids flag drive detection enrichment: to_ids set to true marks an indicator intended for detection, which is why the framework filters on it (EIM1). The

© CYberSynchorny, 2025

threat_level_id ranges from 1 (High) to 4 (Undefined) and analysis from 0 (Initial) through 2 (Completed). The distribution level (0 for the originating organization only, through 3 for all communities, with 4 for a sharing group) and the TLP Tag together evidence that intelligence is shared under controlled access (DS). The enabled and last-fetched state from GET /feeds evidences that the intelligence informing detection is current.

I.6 Kitsune API

Overview. Kitsune provides external threat intelligence, credential-leak detection, and structured incident alerts. The service base is https://home.kitsuneapp.eu. The framework uses it to check monitored identifiers against leak datasets, to pull curated threat feeds, and to read the structured alerts that carry a severity classification. The feeds and credential-leak lookups support incident detection (EIM1), and the severity-classified alerts support incident analysis and handling (EIM2).

Authentication. Kitsune uses a Bearer JWT. The token is sent in the Authorization header as a bearer credential, with the scheme declared as bearer and the format as JWT in the service's security definition.

Request body for lookups. The lookup endpoints share a single request object: a required query string (the identifier or term to search) and an optional next string used to page through results.

```
JSON
POST /api/do/entity/email-domain?datefrom=2026-05-01&dateto=2026-05-19
{
  "query": "amc.example",
  "next": null
}
```

Endpoint groups.

- **Entity lookups**: POST /api/do/entity/{type}, where {type} is one of bank, card, crypto, email-address, email-domain, or ip, with optional datefrom and dateto query parameters. These check a monitored identifier, for example a corporate email domain or an IP, against the underground and leak datasets.
- **Cross-dataset search**: POST /api/do/v2/lookup, a general search across the datasets using the same request body.
- **Stealer-log records**: POST /api/i/bots to search bot and stealer-log records, GET /api/i/bots/{indexfile} to retrieve a record tree, and GET /api/i/bots/{indexfile}/download to download it.
- **Record view**: GET /api/i/view/{bucket}/{systemid} to retrieve a specific record.
- **Live leak lookup**: POST /api/i/3/lookup/live to submit a live lookup and GET /api/i/3/lookup/live/{key} to retrieve its result.
- **Threat feeds**: nine GET endpoints under /api/feeds/, namely ransomware_url, apt_hash, apt_ip, apt_url, botnet_cnc_url, ip_reputation, malicious_hash, malicious_url, and phishing_url, which supply the curated indicators correlated against telemetry.

Alert and incident schema. Structured alerts are returned in a paginated list (count, next, previous, and a results array). Each alert has the following fields:

© CYberSynchorny, 2025

```
JSON
{
  "id": 48213,
  "service_route": "credential-leak",
  "customer_fk": "b1e7c0a2-9f44-4c1a-9d2e-uuid",
  "ks_id": "KS0412",
  "severity": "HIGH",
  "summary": "Corporate credentials exposed in a data-leak dataset",
  "date": "2026-05-19",
  "object_id": "leak-2026-0512",
  "tags": { "source": "underground", "type": "credential" },
  "additional_informations": "...",
  "overview": "...",
  "matched_string": "user@amc.example",
  "is_visualized": false
}
```

Field semantics and control mapping. The severity field is one of INFO, LOW, MEDIUM, HIGH, or CRITICAL, and is the value the EIM2 handling logic prioritizes on. The matched_string is the monitored identifier that triggered the alert, for example a leaked corporate address, which is what the credential-leak check reads. The service_route identifies which Kitsune service produced the alert, tags is a free-form object of classification keys, and date carries the alert timestamp. The feed endpoints and the entity-lookup results provide the detection indicators (EIM1), while the severity-classified alerts drive the incident-handling KCI (EIM2).

I.7 Keycloak Identity Services

Overview. Keycloak (D6.5) is both the framework's authentication provider and a source of identity and access evidence. The framework verifies token claims, queries the Admin REST API for MFA and role evidence, and issues remediation actions. It supports strong authentication (IAM4), role-based access control (IAM1), and the role-assignment aspect of governance (GOV1).

OIDC endpoints (per realm).

- GET /realms/{realm}/.well-known/openid-configuration: discovery of the realm's endpoints.
- POST /realms/{realm}/protocol/openid-connect/token: token issuance, using the Authorization Code flow for user sessions and the Client Credentials grant for the framework's service account.
- GET /realms/{realm}/protocol/openid-connect/certs: the JWKS used to verify token signatures locally.
- GET /realms/{realm}/protocol/openid-connect/userinfo: user claims for an access token.

Token claims. The verified JWT payload carries the identity and authentication context the framework reads:

© CYberSynchorny, 2025

```
JSON
{
  "sub": "f7c3a1d2-8b9e-4c0a-9f11-uuid",
  "preferred_username": "a.skarpeti",
  "acr": "1",
  "amr": ["pwd", "otp"],
  "realm_access": { "roles": ["auditor", "cis_admin"] },
  "resource_access": { "cybergoplus": { "roles": ["compliance-viewer"] } },
  "iss": "https://idp.example/realms/cybersynchrony",
  "aud": "cybergoplus",
  "iat": 1747560600,
  "exp": 1747560900
}
```

The acr (authentication context class reference) and amr (authentication methods references) confirm whether a session used multi-factor authentication, with an amr containing otp indicating a second factor. The realm_access.roles and resource_access.{client}.roles are reconciled against the Role Assignment Matrix. The sub and preferred_username attach evidence to an accountable identity.

Admin REST evidence endpoints. Called with the service-account token, scoped to the view-users and view-events roles:

- GET /admin/realms/{realm}/users: the user population, each with id, username, enabled, and createdTimestamp.
- GET /admin/realms/{realm}/users/{id}/credentials: the credentials per user; the presence of an entry of type otp evidences MFA enrollment.
- GET /admin/realms/{realm}/authentication/flows: the authentication flows, used to confirm that the browser flow requires OTP for privileged groups.
- GET /admin/realms/{realm}/authentication/required-actions: the configured required actions, such as CONFIGURE_TOTP and UPDATE_PASSWORD.
- GET /admin/realms/{realm}/events?type=LOGIN: login events as runtime authentication evidence.
- GET /admin/realms/{realm}/roles and GET /admin/realms/{realm}/users/{id}/role-mappings: the realm roles and a user's mappings, for role reconciliation.

A credentials response and a login event are shaped as follows:

© CYberSynchorny, 2025

```
JSON
[
 { "id": "c1a2", "type": "password", "createdDate": 1700000000000 },
 { "id": "c3b4", "type": "otp", "createdDate": 1710000000000 }
]
```

```
JSON
{
 "time": 1747560600000,
 "type": "LOGIN",
 "realmId": "cybersynchrony",
 "userId": "f7c3a1d2-8b9e-4c0a-9f11-uuid",
 "ipAddress": "100.64.12.20",
 "details": { "username": "a.skarpeti", "auth_method": "openid-connect" }
}
```

Remediation actions. The remediation workflow (REQ-FUNC-08) issues the following writes. MFA enrollment is enforced by setting the CONFIGURE_TOTP required action, a password reset by setting UPDATE_PASSWORD, and active sessions are terminated through the logout endpoint. Role deviations are corrected through the realm role-mapping endpoints.

```
JSON
PUT /admin/realms/{realm}/users/{id}
{ "requiredActions": ["CONFIGURE_TOTP"] }
```

```
JSON
POST /admin/realms/{realm}/users/{id}/logout
```

Control mapping. The acr and amr claims, the otp credential, and the OTP-required browser flow together evidence IAM4; the realm roles and user role-mappings reconciled against the matrix evidence IAM1 and the role-assignment aspect of GOV1; the login events provide runtime authentication evidence.

I.8 CYBERGOPLUS Presentation API

Overview. CYBERGOPLUS exposes a read-only REST API through its API gateway (Section 6.3.3), from which the platform UI and the board-level dashboard fetch the compliance posture. The API returns JSON over HTTPS.

Authentication. Each request carries a Keycloak-issued JWT. The gateway validates the signature against the realm JWKS and scopes the response to the roles asserted in the token, so a caller receives only the compliance data its role permits.

© CYberSynchorny, 2025

Endpoints.

- GET /api/v1/compliance/status: returns the ComplianceStatusJSON object for the summary view.
- GET /api/v1/compliance/findings: returns the open findings, filterable by control, severity, or assessment area.
- GET /api/v1/compliance/kpis: returns the current KPI scores with their methodology metadata.
- GET /api/v1/compliance/report?format={pdf|csv}: returns the exportable board report (REQ-FUNC-09), with Content-Type application/pdf or text/csv.

ComplianceStatusJSON. The summary object carries the overall posture, the per-area breakdown, the KPI set, the per-control status, and the open findings:

```
JSON
{
 "generated_at": "2026-05-19T12:00:00Z",
 "overall_maturity_level": 3,
 "overall_compliance_percentage": 78.4,
 "assessment_areas": [
  { "area": "Governance and Strategy", "maturity_level": 2, "compliance_percentage": 61.0 },
  { "area": "Asset Security",          "maturity_level": 3, "compliance_percentage": 80.0 },
  { "area": "Identity Control",        "maturity_level": 4, "compliance_percentage": 92.0 },
  { "area": "Threat Protection",       "maturity_level": 3, "compliance_percentage": 77.0 },
  { "area": "Vulnerability Status",    "maturity_level": 3, "compliance_percentage": 74.0 },
  { "area": "Resilience",              "maturity_level": 3, "compliance_percentage": 79.0 }
 ],
 "kpis": [
  {
   "id": "KPI 9.2",
   "name": "Industry-specific Standard Adoption",
   "value": null,
   "baseline": 58,
   "target": 65,
   "measurement_window": "project",
   "evidence_source": "use-case survey (D8.3)"
  }
 ],
```

(continued next page)

© CYberSynchorny, 2025

```
JSON
"controls": [
  { "control": "IAM4", "status": "pass", "last_evaluated": "2026-05-19T11:55:00Z", "evidence_source": "Keycloak" },
  { "control": "VM1", "status": "fail", "last_evaluated": "2026-05-19T11:50:00Z", "evidence_source": "Wazuh, SNYK" },
  { "control": "NS4", "status": "no-evidence", "last_evaluated": "2026-05-19T11:40:00Z", "evidence_source": "Suricata" }
 ],
 "open_findings": [
  {
   "finding_id": "F-2026-0512",
   "control": "VM1",
   "severity": "high",
   "asset_id": "agent:002",
   "status": "open",
   "created_at": "2026-05-19T11:50:00Z",
   "due_date": "2026-05-26"
  }
 ]
}
```

The numeric values above are illustrative of the response structure and are not measured pilot results. The KPI 9.2 entry shows the real defined baseline and target with a null value, indicating an indicator whose measurement is pending the survey reported in D8.3.

Field semantics. The overall_maturity_level is an integer from 1 to 5 and overall_compliance_percentage a value from 0 to 100. Each assessment_areas entry reports one of the six domains of Section 5.2 with its own maturity and percentage. Each kpis entry carries the methodology metadata of Section 5.3.4, namely the baseline, target, measurement_window, and evidence_source, with value null until measured. Each controls entry reports a control status of pass, fail, or no-evidence, the last of which encodes the fail-safe rule of Section 10.3, in which absent evidence is surfaced rather than assumed to pass. Each open_findings entry carries the finding identifier, the affected control and asset, the severity, the status, and the due date used by the remediation workflow.

© CYberSynchorny, 2025

Annex II: Standard Adoption Baseline Evidence (KPI 9.2)

To ensure that every reported KPI is interpretable and verifiable, this version specifies a common measurement methodology. Each KPI is documented with five attributes: a baseline (the as-is value or condition from which change is measured), a target population (the entities or assets to which it applies), a measurement window (the period and cadence of measurement), an evidence source (the artifact or dataset from which the value is computed), and a statement of confidence and limitations. KPIs are further classified as preparatory indicators, which evidence that the capability to measure an outcome has been put in place, or outcome indicators, which report a measured result. The operational KPIs in Table 9 follow this template, with their values computed from the KCIs described above. This section specifies in full the outcome KPI for which D6.8 is the designated methodological home, KPI 9.2, while the measured result is reported in D8.3.

Cross-Sector Maturity Picture

The sectors are not starting from a common baseline. Those that already had dedicated supervisory frameworks, mature incident-sharing communities, or operational safety cultures entered the NIS2 era with higher adoption capacity than sectors newly brought into scope or composed of smaller, dispersed operators. The comparative picture from ENISA's cross-sector work and threat-landscape evidence is summarized below.

Table 13: Comparative Standard Adoption by Critical Sector

Sector	Adoption picture	Key measurable indicators	Principal standards and rules
Energy	Uneven; electricity is top-tier and highly mature, gas moderate, oil, hydrogen and district heating lag	No harmonized EU compliance percentage; ENISA marks electricity as both highly mature and highly critical	NIS2; electricity cybersecurity network code (EU) 2024/1366; ENISA NIS360
Transport	Moderate overall, aviation leads; rail and maritime mid-tier; road weaker	ENISA 2024 threat landscape attributed about 11% of observed incidents to transport	NIS2; aviation Regulations (EU) 2023/203 and (EU) 2022/1645; EASA and ERA guidance
Health (UC1)	Upper-moderate but highly heterogeneous; large providers and manufacturers stronger; legacy systems and medical devices a persistent drag	ENISA recorded 487 health incidents in 2024 (about 45% ransomware, 28% data breaches); most-affected sector in CIRAS from 2020 to 2023	NIS2; ENISA health threat-landscape work; hospital cyber-resilience guidance
Finance	High and rising, driven by DORA; banking more mature than financial market infrastructures	ENISA attributed about 9% of 2024 incidents to finance; 12% of significant NIS incidents in 2023	NIS2; DORA (sector-specific from 17 January 2025); ECB and ESA guidance
Digital infrastructure and telecoms	Among the strongest overall but not uniform; telecoms high, cloud, data centres and trust services moderate-high	ENISA counted 188 major telecom incidents and 44 trust-service incidents in 2024; about half of EU countries use the EECC technical guidelines	EECC; eIDAS and EUDI; Implementing Regulation (EU) 2024/2690; ENISA 2025 technical guidance

© CYberSynchorny, 2025

Water	Still weak, especially wastewater; drinking water moderate	ENISA: drinking water was about 1% of 2023 CIRAS incidents (44% malicious); no EU-level sector cyber drill recorded	NIS2; national guidance where available
Public administration (UC2)	One of the clearest lagging sectors; placed in the risk zone	80% report leadership approval of cyber-risk controls, 60% report supply-chain risk policies, 76% report no information-sharing participation; 19% of 2024 threat-landscape incidents	NIS2; national templates and self-assessment tools

The academia case (UC3) is not a core NIS2 sector in the available evidence, so no sectoral adoption proxy exists for it; its baseline is established from the project's own assessment.

Entity-Size Adoption Gap

The clearest cross-cutting pattern is entity size. Eurostat's 2024 enterprise data shows maturity rising sharply with size: strong-password authentication is used by 96.78% of large enterprises against 82.03% of small enterprises, and formal ICT-security documentation was defined or reviewed in the previous twelve months by 57.87% of large enterprises against 17.98% of small enterprises. ENISA's sector work reaches the same conclusion qualitatively, with smaller and newly scoped operators facing staffing and budget deficits. This size effect is material for the use cases, whose adoption baselines depend heavily on the size and resourcing of the participating entities.

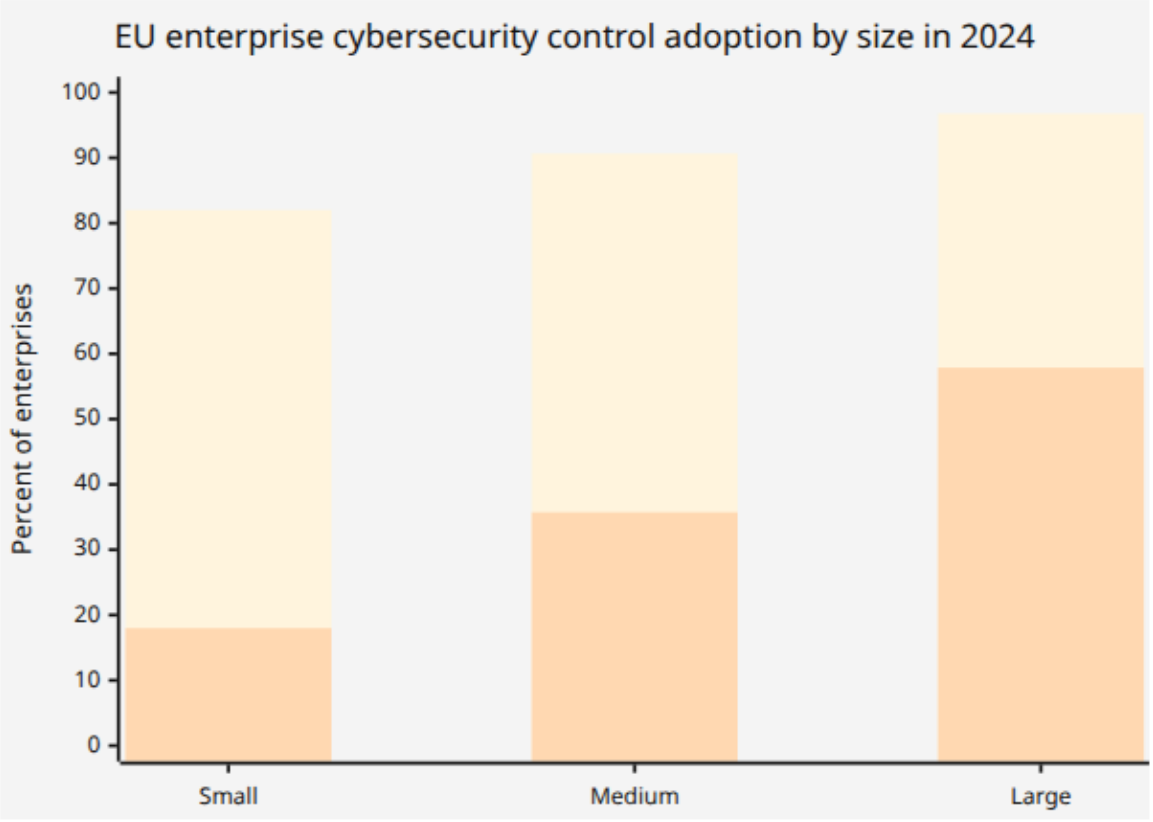

Figure 5: EU enterprise cybersecurity control adoption by size, 2024 (Eurostat)

© CYberSynchorny, 2025

Legal Adoption and Enforcement Context

Legal adoption has been uneven. As of the Commission's transposition state of play dated 1 July 2025, only eight Member States were recorded as having fully notified transposition, while nineteen had received a reasoned opinion on 7 May 2025 for failure to notify. Enforcement escalated from formal notices issued to 23 Member States in late 2024 to reasoned opinions to 19 Member States in May 2025. Rising incident-report volumes over this period are best read as improved reporting culture and visibility rather than only as deteriorating conditions. The implementation timeline and the enforcement trajectory are shown below.

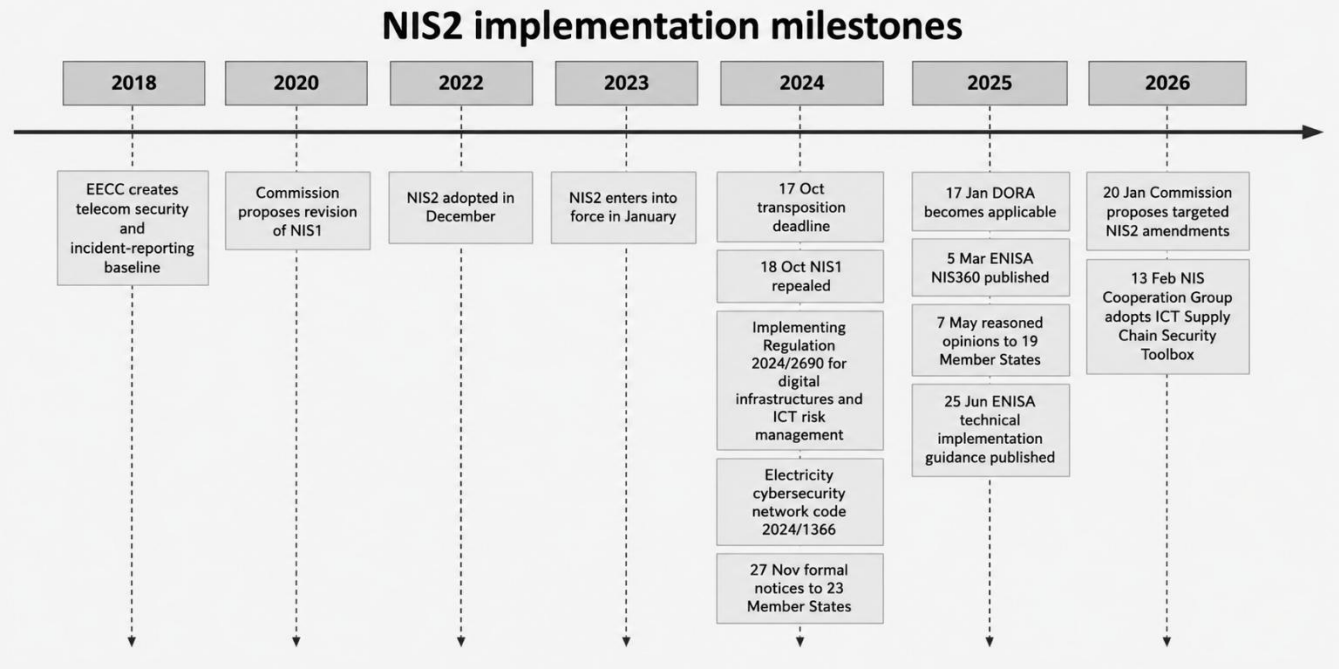

Figure 6: NIS2 implementation milestones, 2018 to 2026

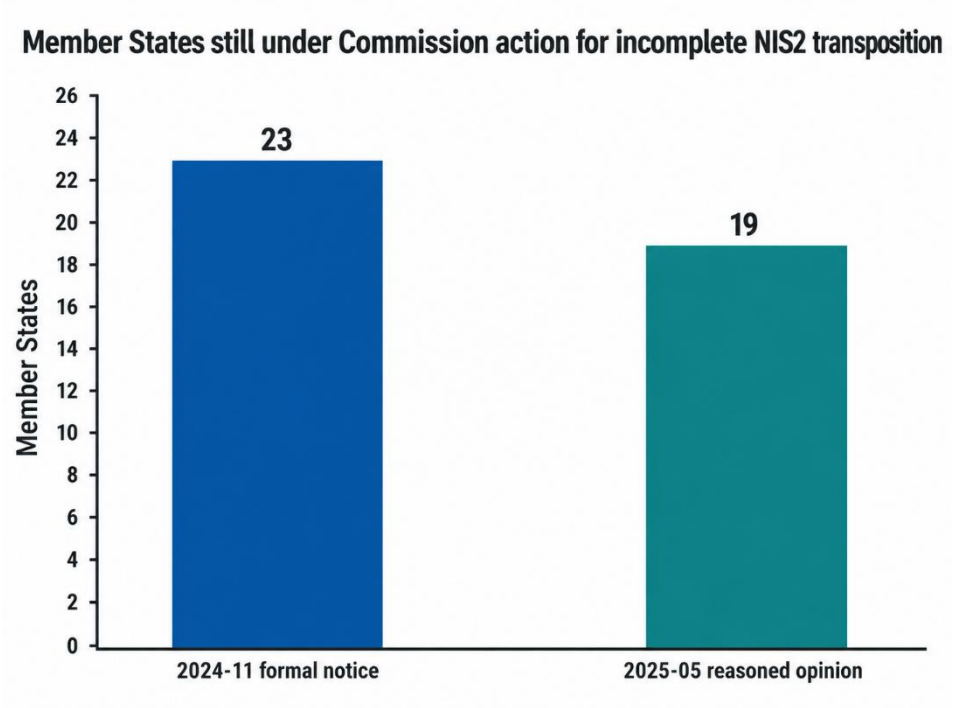

Figure 7: Member States under Commission action for incomplete NIS2 transposition

© CYberSynchorny, 2025

ENISA's 2024 baseline survey found newly covered NIS2 sectors markedly weaker than long-regulated sectors in information sharing, preparedness, and supply-chain control, reporting over 60% non-participation in information-sharing initiatives among newly covered sectors and noting that around 20% of those entities still implicitly trusted their supply chains. This is among the best available Union-wide indicators that supply-chain and standard adoption currently lag behind legal transposition.

Consolidated Baseline

For the purpose of KPI 9.2, the proxy indicators above are consolidated into a single baseline adoption rate of approximately 58%. This figure is anchored on the Eurostat 2024 measure of formal ICT-security documentation maintained by large enterprises (57.87%), the size class of the use-case entities, and is consistent with ENISA's maturity placement of the use-case sectors. It is a triangulated proxy rather than a direct compliance count, and it will be superseded by the project's own survey of the use-case participants, whose result is reported in D8.3 against the 65% target.

Sources

- European Union Agency for Cybersecurity (ENISA), NIS360 sector maturity and criticality assessment, 2025.
- ENISA, NIS Investments report series.
- ENISA, Threat Landscape 2024.
- ENISA, 2024 NIS baseline survey of newly covered sectors.
- Eurostat, ICT usage and security in enterprises, 2024.
- European Commission, NIS2 transposition state of play, as of 1 July 2025.
- ENISA, NIS2 technical implementation guidance, 2025.
- Sector instruments: Electricity cybersecurity network code (EU) 2024/1366; Aviation Regulations (EU) 2023/203 and (EU) 2022/1645; European Electronic Communications Code; eIDAS; Implementing Regulation (EU) 2024/2690; Digital Operational Resilience Act (DORA).

© CYberSynchorny, 2025