



AI-DRIVEN CYBERSECURITY FOR FINANCIAL SERVICE PROVIDERS

D2.1 – FINANCIAL SECTOR CYBERSECURITY REQUIREMENTS ANALYSIS

Lead Beneficiary	UBITECH
Work Package Ref.	WP2-LLM-Powered Reference Architecture, Specification and Systems Integration
Task Ref.	T2.1 Cybersecurity Requirements and Reference Blueprints
Deliverable Title	Financial Sector Cybersecurity Requirements Analysis
Due Date	30/06/2026
Delivered Date	30.06/2026
Revision Number	1.0
Dissemination Level	PU - Public
Type	R — Document, report
Document Status	Final
Review Status	Reviewed
Document Acceptance	Acceptable
EC Project Officer	Rodica Morariu

DIGITAL-ECCC-2024-DEPLOY-CYBER-07-KEYTECH



Co-funded by
the European Union



ECCC 
EUROPEAN CYBERSECURITY
COMPETENCE CENTRE

The project funded under Grant Agreement No 101249596 is supported by the European Cybersecurity Competence Centre

DISCLAIMER

The sole responsibility for the content of this publication lies with the authors. It does not necessarily reflect the opinion of the European Union. Neither the EASME nor the European Commission is responsible for any use that may be made of the information contained therein.

COPYRIGHT MESSAGE

This report, if not confidential, is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0); a copy is available here: <https://creativecommons.org/licenses/by/4.0/>. You are free to share (copy and redistribute the material in any medium or format) and adapt (remix, transform, and build upon the material for any purpose, even commercially) under the following terms: (i) attribution (you must give appropriate credit, provide a link to the license, and indicate if changes were made; you may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use); (ii) no additional restrictions (you may not apply legal terms or technological measures that legally restrict others from doing anything the license permits).

CONTRIBUTING PARTNERS

Partner Acronym	Role ¹	Name Surname ²
UBI	Editor, Contributor	Nikos Papageorgopoulos, Giannis Ledakis
JRC	Contributor, Internal Reviewer	Konstantina Tripodi
CXB	Contributor, Internal Reviewer	Ramon Martín de Pozuelo, Martí Fabregat
FAIR	Contributor	Fernando Ferreira
HUA	Contributor	George Kousiouris, Panagiotis Rizomiliotis, Christos Diou
INNOV	Contributor	George Fatouros
IQB	Contributor	Cosmina Stalidi
KM3	Contributor	Kostas Metaxas, Konstantinos Ilias
K3Y	Contributor	Evangelia Pantraki, Nikolaos Ntampakis
QUBO	Contributor, Internal Reviewer	Ernstjan de Gooyert
OTE	Contributor	Achilleas Marinakis, Christos Gizelis
UPRC	Contributor	Pepi Paraskevopoulou, Konstantinos Mavrogiorgos, Giorgos Makridis

REVISION HISTORY

Version	Date	Partner(s)	Description
0.1	30-03-2026	UBI	Creation of deliverable document and initial Table of Contents.
0.2	16-04-2026	UBI, JRC	Drafting of Section 1 and contribution to pilot-related content, including introduction, scope, relation to work packages, deliverable structure and JRC pilot input.
0.3	12-05-2026	UBI, IQB, UPRC, QUBO, HUA	Drafting and consolidation of Section 2 and Section 3, including methodology, requirements elicitation approach, financial-sector threat landscape, AI/GenAI threats, supply chain attacks, post-quantum cryptography and threat intelligence landscape.
0.4	18-05-2026	UBI, UPRC, INNOV, IQB, JRC, KM3, HUA, FAIR, QUBO, CXB, K3Y	Development of section 5 and requirements catalogue structure, functional and non-functional requirements matrices, prioritisation and user stories structure.
0.5	05-06-2026	CXB, FAIR, OTE, KM3	Partner contributions to pilot/stakeholder requirements and regulatory and standards framework
0.6	19-06-2026	JRC, CXB	Internal review phase
0.7	26-06-2026	UBI, IQB, FAIR	Changes based on the review
1.0	30-06-2026	UBI	Final quality check and submission

LIST OF ABBREVIATIONS

Acronym	Definition
4AMLD	Fourth Anti-Money Laundering Directive
AI	Artificial Intelligence
AML	Anti-Money Laundering
API	Application Programming Interface
ASVS	Application Security Verification Standard
ATT&CK	Adversarial Tactics, Techniques and Common Knowledge
CI	Critical Infrastructure
CI/CD	Continuous Integration / Continuous Delivery
CRA	Central Router Agent
CSF	Cybersecurity Framework
CTI	Cyber Threat Intelligence
DIOSCURI	Digital Twin-based cybersecurity and orchestration platform
DORA	Digital Operational Resilience Act
DPIA	Data Protection Impact Assessment
EBA	European Banking Authority
ECB	European Central Bank
ECSMAF	ENISA Cybersecurity Market Analysis Framework
EIOPA	European Insurance and Occupational Pensions Authority
ENISA	European Union Agency for Cybersecurity
EU	European Union
FS-ISAC	Financial Services Information Sharing and Analysis Center
GDPR	General Data Protection Regulation
GenAI	Generative Artificial Intelligence
HFT	High-Frequency Trading
HITL	Human-in-the-Loop
ICT	Information and Communication Technology
ISO/IEC	International Organization for Standardization / International Electrotechnical Commission
LLM	Large Language Model
MCP	Model Context Protocol
MITRE	MITRE Corporation
ML-KEM	Module-Lattice-Based Key-Encapsulation Mechanism
NIS2	Network and Information Security Directive 2
NIST	National Institute of Standards and Technology
OASIS	Organization for the Advancement of Structured Information Standards
OLISTIC	Risk assessment and management engine
OWASP	Open Worldwide Application Security Project
PCI-DSS	Payment Card Industry Data Security Standard
PQC	Post-Quantum Cryptography
PSD2	Revised Payment Services Directive

Q2FA	Quantum Two-Factor Authentication
RAG	Retrieval-Augmented Generation
SAI	Securing Artificial Intelligence
SBOM	Software Bill of Materials
SCA	Strong Customer Authentication
SIEM	Security Information and Event Management
SLH-DSA	Stateless Hash-Based Digital Signature Algorithm
SKM	Security Knowledge Model
SOC	Security Operations Centre
STIX	Structured Threat Information Expression
TAXII	Trusted Automated Exchange of Intelligence Information
TIBER-EU	Threat Intelligence-Based Ethical Red Teaming – European Framework
TLPT	Threat-Led Penetration Testing
TLS	Transport Layer Security
UC	Use Case
UEBA	User and Entity Behaviour Analytics
XDR	Extended Detection and Response
XENON	AI-driven penetration testing and red-teaming platform

EXECUTIVE SUMMARY

Deliverable D2.1, Financial Sector Cybersecurity Requirements Analysis, presents the initial cybersecurity requirements baseline for the CyberAid project. The deliverable is developed in the context of WP2 and provides a structured analysis of the needs, challenges and requirements that will guide the design, development, integration and validation of the CyberAid framework.

The work reported in this deliverable is based on inputs from the CyberAid Grant Agreement, partner contributions, technical work packages, pilot needs, regulatory and standards analysis, and the financial-sector threat landscape. The deliverable considers the specific characteristics of financial critical infrastructures, including their dependency on digital services, cloud environments, APIs, third-party providers, transaction systems, trading environments and AI-enabled processes.

The methodology followed in D2.1 includes the elicitation, review, classification, prioritisation and validation of requirements. Requirements were collected from project partners and then structured according to their source, type, priority, related work package, CyberAid capability area and validation means. The MoSCoW method was used to support prioritisation, while user stories were introduced as a way to connect requirements with practical usage scenarios and pilot validation activities.

The deliverable also analyses the regulatory and standards framework relevant to CyberAid, including DORA, NIS2, GDPR, PSD2, 4AMLD and the EU AI Act, as well as standards and frameworks such as ISO/IEC 27001, ISO/IEC 27035, NIST CSF, TIBER-EU, OWASP, MITRE ATT&CK and STIX/TAXII. This analysis supports the alignment of the requirements catalogue with operational resilience, incident reporting, privacy, AI governance, secure development and threat intelligence needs.

The resulting requirements catalogue is organised according to the main CyberAid capability areas: CyberAid-BLUEPRINTS, CyberAid-PROACTIVE, CyberAid-MONITOR and CyberAid-REPORT. Functional and non-functional requirements are presented separately, while pilot and stakeholder requirements, user stories, prioritisation and traceability information provide additional views of the same baseline.

Overall, D2.1 establishes the first consolidated requirements baseline of CyberAid. This baseline will support the reference architecture, cybersecurity blueprints, technical specifications, component development, integration work, pilot validation and later evaluation activities of the project.

1.	Introduction	9
1.1	Scope of Deliverable.....	9
1.2	Relation to other Work Packages and Work Plan.....	10
1.3	Deliverable Structure.....	10
2.	Methodology.....	12
2.1	Requirements Elicitation Approach.....	12
2.2	Sources of Requirements	14
2.3	Requirements Classification and Prioritisation	15
2.4	Stakeholder Engagement and Validation Process	16
3.	Financial Sector Context and Threat Landscape	18
3.1	Overview of the Financial Sector as Critical Infrastructure	18
3.2	Financial-Sector Attack Surface and Key Attack Vectors.....	19
3.3	Emerging Threats in Finance	19
3.3.1	GDPR AI/GENAI related vulnerabilities and threats.....	21
3.3.2	Supply Chain Attacks	22
3.3.3	Post Quantum Cryptography	24
3.4	Threat Intelligence Landscape	26
4.	Regulatory and Standards Framework.....	28
4.1	EU Regulatory Framework Overview.....	28
4.1.1	DORA (Digital Operational Resilience Act).....	28
4.1.2	NIS2 Directive (Network and Information Security Directive 2)	29
4.1.3	GDPR (General Data Protection Regulation).....	30
4.1.4	PSD2 and Strong Customer Authentication	31
4.1.5	4AMLD (4th Anti-Money Laundering Directive)	31
4.1.6	EU AI Act.....	32
4.2	Relevant Standards and Frameworks	33
4.2.1	ISO/IEC 27001 and 27035.....	34
4.2.2	NIST Cybersecurity Framework (CSF)	34
4.2.3	TIBER-EU Framework.....	34
4.2.4	OWASP Financial API Security Guidelines.....	35
4.2.5	MITRE ATT&CK and STIX/TAXII	35
5.	CyberAid Requirements Catalogue	37
5.1	Overview of CyberAid Requirements.....	37
5.2	Functional Requirements	40
5.2.1	Blueprint Requirements.....	40
5.2.2	Monitor Requirements	41
5.2.3	Proactive Requirements.....	44
5.2.4	Report Requirements.....	46

5.3	Non-Functional Requirements	47
5.3.1	Blueprint Requirements.....	47
5.3.2	Monitor Requirements	48
5.3.3	Proactive Requirements.....	48
5.3.4	Report Requirements.....	49
5.4	Pilot Stakeholder Requirements.....	50
5.4.1	Pilot #1 – Client Impersonation Detection	50
5.4.2	Pilot #2 – Payment Processing / iSOC	51
5.4.3	Pilot #3 –Integrated Security Operations	53
5.4.4	Pilot #4 – Financial Market and Trading Systems	55
5.4.5	Cross-Pilot Requirements Contributions.....	58
5.5	User Stories	59
5.5.1	User Stories Definition Approach	59
5.5.2	User Stories Matrix	59
6.	Conclusions And Next Steps	63

LIST OF FIGURES

Figure 1: CyberAid Requirements Engineering Workflow	12
Figure 2: Sources of Requirements	15
Figure 3: Stakeholder Validation and Feedback Loop	17

LIST OF TABLES

Table 1: Requirements Collection Template	13
Table 2: MoSCoW Prioritisation Template	16
Table 3: User Stories Template	16
Table 4: Requirement Validation Status	17
Table 5: Mapping of EU regulatory requirements to CyberAid capabilities	33
Table 6: Mapping of relevant standards and frameworks to CyberAid components	36
Table 7: Requirement Catalogue.....	39
Table 8: Blueprint Functional Requirements.....	41
Table 9: Monitor Functional Requirements	43
Table 10: Proactive Functional Requirements	45
Table 11: Report Functional Requirements	47
Table 12: Blueprint Non-Functional Requirements.....	48
Table 13: Monitor Non-Functional Requirements.....	48
Table 14: Proactive Non-Functional Requirements.....	49
Table 15: Report Non-Functional Requirements.....	49
Table 16: Cross-pilot contribution to CyberAid requirements	59
Table 17: Initial CyberAid User Stories Matrix.....	62

1. INTRODUCTION

The increasing digitalisation of financial services, combined with the growing sophistication of cyber threats have significantly increased the attack surface of financial institutions and critical financial infrastructures. Financial organisations are constantly at risk of threats targeting operational continuity, sensitive financial data, digital assets, payment systems, and interconnected supply chains. At the same time, the evolving European regulatory landscape, including frameworks such as the Digital Operational Resilience Act (DORA), the NIS2 Directive, and the EU AI Act, introduces additional requirements related to cybersecurity resilience, risk management, incident handling, and trustworthy AI adoption.

Within this context, the CyberAId project aims to develop an AI-driven cybersecurity framework that can enhance cyber defence, situational awareness, threat detection, and incident response capabilities for the financial sector. The project focuses on improving the resilience and trustworthiness of financial infrastructures through the integration of advanced cybersecurity technologies, intelligent analytics, automation mechanisms, and collaborative threat intelligence approaches.

Deliverable D2.1 constitutes one of the initial outcomes of WP2 and provides the baseline requirements analysis for the CyberAId framework. The purpose of this deliverable is to identify, analyse, structure, and prioritise the cybersecurity requirements relevant to the financial sector and the project use cases. The deliverable consolidates technical, operational, regulatory, and stakeholder-driven requirements that will guide the specification, design, development, and integration activities performed in subsequent work packages.

The work presented in this deliverable is based on the analysis of the financial sector threat landscape, existing cybersecurity challenges, relevant regulatory and standardisation frameworks, project objectives defined in the Grant Agreement, and inputs collected from project partners and pilot stakeholders. The identified requirements establish the foundation for the definition of the CyberAId architecture and ensure alignment between technological developments and the operational needs of the targeted financial environments.

The remainder of this deliverable is structured as follows. Section 2 presents the methodology followed for requirements elicitation, analysis, validation, and prioritisation. Section 3 analyses the financial sector context and the cybersecurity challenges to be addressed by the project. Section 4 presents the regulatory and standardisation landscape relevant to the financial domain. Section 5 introduces the consolidated requirements catalogue, including functional and non-functional requirements, prioritisation, and traceability aspects. Finally, Section 6 summarises the main conclusions and lays out the next steps of the project.

1.1 Scope of Deliverable

This deliverable defines the initial cybersecurity requirements baseline for the CyberAId project, focusing on the needs, challenges, and operational conditions of the financial sector. The scope of this deliverable will be to identify, analyse, classify and prioritise the technical, operational, security, regulatory, and user-oriented requirements that will guide the development of the CyberAId framework and its associated components.

The work performed within the context of the deliverable addresses the cybersecurity challenges faced by financial organisations in highly digitalised and interconnected environments, taking into consideration both current and emerging threats, as well as the increasing adoption of AI-enabled technologies within financial infrastructures. Particular emphasis is placed on requirements associated with cyber threat detection, incident response, operational resilience, data protection, secure information exchange, AI trustworthiness, and compliance with the relevant European regulatory frameworks.

The requirements elicitation process is based on inputs from multiple sources, including the CyberAId Grant Agreement, project use cases, partner expertise, pilot stakeholders' needs, cybersecurity standards, and

applicable European regulations and policies. The resulting requirements catalogue establishes the foundation for the design and implementation activities that will be carried out in the subsequent technical work packages.

Furthermore, this deliverable provides an initial structure for requirements traceability and prioritisation, ensuring that future development, integration, validation, and demonstration activities remain aligned with the project objectives and the operational expectations of the targeted financial sector stakeholders.

1.2 Relation to other Work Packages and Work Plan

The work presented in Deliverable D2.1 is closely connected with the overall CyberAId work plan and serves as a foundational input for several subsequent work packages and technical activities. The identified requirements establish the baseline for the specification, development, integration, validation, and demonstration of the CyberAId framework and its components.

More specifically, the results of this deliverable are directly supporting the technical design and architecture definition activities that will be performed within the technology-oriented work packages. The prioritised requirements catalogue provides guidance for the development of AI-driven cybersecurity mechanisms, data processing pipelines, threat intelligence functionalities, monitoring services, and decision-support capabilities foreseen within the project.

The requirements identified in this deliverable will also contribute to the definition of validation and evaluation criteria within the pilot and demonstration activities. The mapping between stakeholder needs, operational scenarios, and technical functionalities enables the definition of measurable objectives and facilitates the assessment of the effectiveness, usability, resilience, and trustworthiness of the developed solutions in realistic financial-sector environments.

Moreover, the deliverable contributes to ensuring alignment between CyberAId technological developments and the applicable regulatory, ethical, and legal requirements associated with cybersecurity, operational resilience, privacy preservation, and trustworthy AI adoption within the European financial ecosystem.

The outputs of D2.1 will be continuously refined and updated throughout the project lifecycle as additional technical insights, stakeholder feedback, pilot outcomes, and evolving regulatory requirements become available. Consequently, this deliverable acts as an initial reference point for maintaining consistency and traceability across the entire CyberAId development and validation process.

In this sense, D2.1 acts as an initial requirement baseline for several subsequent activities of the project. The requirements defined in this deliverable are expected to feed the WP2 outputs, including the reference architecture, blueprints, security knowledge models, LLM orchestration design and integration work, but also the technical development, pilot validation, evaluation and demonstration activities performed in the other work packages.

1.3 Deliverable Structure

This deliverable is organised into several sections covering the methodology, analysis activities, and resulting cybersecurity requirements relevant to the CyberAId project objectives.

Section 2 presents the methodology followed for requirements elicitation, analysis, validation, and prioritisation. It describes the different information sources considered during the requirements engineering process, including stakeholder inputs, regulatory frameworks, standards, use cases, and cybersecurity best practices.

Section 3 analyses the financial-sector context and the cybersecurity challenges addressed within the scope of CyberAid. The section discusses the evolving threat landscape, operational constraints, sector-specific cybersecurity needs, and the role of AI-driven cybersecurity solutions within financial environments.

Section 4 presents the regulatory, legal, and standardisation landscape relevant to the project. This includes the analysis of European cybersecurity and resilience frameworks, AI-related regulations, privacy and data protection considerations, and applicable standards influencing the design and operation of the CyberAid framework.

Section 5 introduces the consolidated cybersecurity requirements catalogue, including functional and non-functional requirements, pilot and stakeholder requirements, user stories, prioritisation and traceability aspects.

Finally, Section 6 summarises the main conclusions of the deliverable and outlines the next steps related to the evolution, refinement, and implementation of the identified requirements within the subsequent phases of the project.

2. METHODOLOGY

This section presents the methodology followed for the collection, analysis, classification, prioritisation and validation of the CyberAid cybersecurity requirements. The methodology was mainly connected with the work performed under T2.1, where the initial requirements baseline was prepared as input for the reference architecture, cybersecurity blueprints, technical specifications and later integration activities.

The approach followed in this deliverable was based on different types of inputs. These included the CyberAid Grant Agreement, technical work package needs, standards, legislation, pilot needs, financial-sector cybersecurity challenges and partners' expertise. This was considered necessary because CyberAid was not only developing individual technical components, but also an integrated AI-driven cybersecurity framework for financial infrastructures and critical information environments.

The methodology also considered selected concepts from the ENISA Cybersecurity Market Analysis Framework (ECSMAF), mainly as a supporting reference for structuring the analysis of the target domain. In particular, the ECSMAF concepts of scope definition, stakeholder identification, asset and threat analysis, security requirement identification and validation were considered relevant for the D2.1 requirements engineering process. However, the ECSMAF was not applied as a full market analysis methodology.

The requirements work was performed iteratively. First, partners collected and documented candidate requirements from technology, standards, legislation and pilot needs. Then, the collected requirements were reviewed and discussed within WP2. After this, the requirements were classified and prioritised using the MoSCoW method. Where useful, requirements were also transformed into user stories, so that they could be expressed from the point of view of the relevant user, analyst, pilot stakeholder or technical actor. The validated requirements and user stories were then mapped, where possible, pilot scenarios, project tasks and CyberAid technical components.

The methodology applied in this deliverable produced the initial CyberAid requirements baseline. This baseline will remain open to updates during the project, as more detailed input becomes available from the architecture work, pilot definition, component development, integration activities and validation exercises. The overall requirements engineering workflow followed is presented in Figure 1.




Figure 1: CyberAid Requirements Engineering Workflow

2.1 Requirements Elicitation Approach

The requirements elicitation approach aimed to collect the needs and constraints that should be taken into account during the design and implementation of the CyberAid framework. The process within the context of T2.1, with input expected from technical partners, pilot partners and partners responsible for legal, regulatory and domain-specific aspects.

Requirements were collected in a shared working document. Each partner contributed requirements based on its role in the project, technical background, expected component, pilot environment or regulatory expertise. The collected requirements were then reviewed in order to identify missing information, overlaps, unclear wording, possible duplicates and dependencies with other requirements.

The requirements were not treated as final at the moment of collection. They were reviewed, refined and discussed with the relevant partners, mainly through WP2 interactions and dedicated requirements discussions. This was important because several requirements needed further clarification once the architecture and the pilot scenarios became more mature.

The elicitation process followed a practical approach. Each requirement was expected to describe a need that was relevant for CyberAId and should be traceable to one or more sources. Such sources included regulations, standards, pilot needs, technical tasks, project objectives, cybersecurity challenges or partner inputs. This helped to avoid a generic list of requirements and supported a more traceable and project-specific requirements catalogue.

The requirements collected in D2.1 were expected to support the later work of WP2, WP3, WP4, WP5 and WP6. In particular, they were intended to be used as input for the reference architecture and blueprints, the definition of technical specifications, the development of cybersecurity components, the pilot validation activities and the integration work.

Requirements were initially documented in a consolidated requirements repository containing all relevant metadata for traceability, prioritisation and validation purposes. Since the complete requirements table contains a large amount of information, the requirements may later be presented in separate matrices throughout the deliverable depending on the analysis needs. For example, one matrix may include only the requirement identifiers and titles, while other matrices may focus on prioritisation, validation means, source mapping, WP/task mapping or requirement classification. This approach improves readability while maintaining traceability to the original consolidated requirements repository. The requirements collection template used to structure this information is presented in Table 1.

Field	Description
Req ID	Final unique identifier of the requirement
Source Type	Type of source from which the requirement originates (e.g., regulation, standard, pilot input, partner input, GA, technical task)
Source Reference	Specific reference to the originating source
Requirement Title	Short descriptive title of the requirement
Requirement Description	Detailed description of the requirement and the need to be addressed
Validation Means	Method used to validate the requirement (e.g., testing, pilot validation, review, demonstration, workshop)
Functional / Non-Functional	Classification of the requirement
MoSCoW Priority	Must Have, Should Have, Could Have or Won't Have
Contributor of Requirement	Partner or stakeholder contributing the requirement
Related WP or Task	WP or task linked to the requirement
Notes	Additional comments, assumptions or dependencies
Status	Draft, under review, accepted, revised, deferred or rejected

Table 1: Requirements Collection Template

2.2 Sources of Requirements

The requirements analysed in this deliverable were derived from different sources. This was necessary because CyberAId covers several areas at the same time, including financial-sector cybersecurity, AI-based orchestration, monitoring, proactive defence, incident response, regulatory reporting, privacy-preserving analytics and pilot validation.

The first source was the CyberAId Grant Agreement. The Grant Agreement defined the project objectives, technical direction, work packages, deliverables and expected outcomes. These elements provided the initial basis for defining what the CyberAId framework should support. The Grant Agreement was also used to keep the requirements aligned with the expected project outputs and the role of WP2 in the overall work plan.

The second source was the technical work of the project. Inputs from WP2, WP3, WP4 and WP5 were used in order to identify requirements related to architecture, cybersecurity tools, CTI integration, monitoring, vulnerability management, privacy-preserving analytics, LLM-based orchestration, reporting and integration. These requirements were important because they connected the project objectives with the technical components that would be designed and implemented.

The third source was the regulatory and standards landscape. Requirements were mapped, where possible, to relevant frameworks such as DORA, NIS2, GDPR, PSD2, the EU AI Act, ISO/IEC 27001, ISO/IEC 27035, the NIST Cybersecurity Framework and other relevant cybersecurity standards and good practices. This mapping was needed in order to ensure that the CyberAId requirements also reflected compliance needs and financial-sector obligations.

The fourth source was the financial-sector threat landscape and the attack surfaces of financial infrastructures. This included needs related to transaction monitoring, API security, log analysis, authentication monitoring, access control, secure data exchange, privacy-preserving analytics, insider threat detection, regulatory reporting and incident handling. These inputs were used to ensure that the requirements related to real security problems faced by financial organisations.

The fifth source was the input of the project partners and pilot stakeholders. Technical partners provided requirements related to the tools and services that they would develop or integrate. Pilot partners provided operational requirements based on their environments and needs. These inputs were important because they helped to verify whether the requirements were meaningful in practical scenarios and not only from a technical perspective.

The ENISA Cybersecurity Market Analysis Framework was used as an additional methodological reference. It supported the identification of the target scope, infrastructure, assets, threats, stakeholders and security requirements. In CyberAId, this was used in a simplified way, mainly to structure the requirements analysis and to avoid missing important dimensions of the financial cybersecurity domain.

The main sources that contributed to the CyberAId requirements catalogue are summarised in Figure 2. The figure shows how the Grant Agreement, technical work packages, standards and regulations, the financial-sector threat landscape, pilot stakeholders and the ENISA ECSMAF reference framework were used as inputs to the requirements baseline.




Figure 2: Sources of Requirements

Industry reports on cloud and AI adoption in financial services were also considered as contextual input, mainly to understand the operational environment in which CyberAid may be deployed. These sources highlight the increasing use of hybrid cloud services, AI agents, third-party providers, non-human identities and data-intensive AI workflows in financial organisations[38]. This input supports the inclusion of requirements related to cloud security, AI monitoring, identity governance, data leakage prevention and oversight of autonomous or semi-autonomous AI functions.

2.3 Requirements Classification and Prioritisation

After collection, the requirements were classified so that they could be reviewed and used more easily by the project partners. The classification also helped to identify overlaps, dependencies and possible gaps in the requirements baseline.

The first classification separated the requirements into functional and non-functional requirements. Functional requirements described what the CyberAid framework or its components should do. These included, for example, threat monitoring, vulnerability assessment, CTI integration, incident response, automated reporting, LLM-supported analysis or privacy-preserving data processing. Non-functional requirements described how the framework should behave, for example in relation to performance, scalability, interoperability, explainability, availability, data protection and secure deployment.

A second classification linked each requirement with the relevant CyberAid area or capability. This included CyberAid-PROACTIVE, CyberAid-MONITOR, CyberAid-REPORT, CyberAid-BLUEPRINTS, the LLM orchestration layer, CTI models, privacy-preserving analytics, quantum-resistant cryptography, learning and training services, or pilot environments. This classification is expected to support the preparation of the traceability matrix.

A third classification concerned the source of the requirement. Each requirement was linked, where possible, to one or more sources such as regulation, standard, project objective, pilot need, technical partner input or threat landscape analysis. This made the catalogue easier to justify and maintain during the project.

For prioritisation, the MoSCoW method was used. Requirements were marked as Must Have, Should Have, Could Have or Won't Have for the current project scope. This method was considered simple enough to be used by all partners and allowed the consortium to distinguish between essential requirements and lower-priority ones.

The prioritisation took into account several aspects, including regulatory relevance, pilot relevance, technical feasibility, dependency with other requirements, expected value for CyberAid, and importance for the reference architecture and integration work. Requirements that were essential for the project architecture, the pilots, or the regulatory alignment were expected to receive higher priority. Requirements

that were useful but not critical could be marked as Should or Could Have. The prioritisation categories used in the MoSCoW assessment are defined in Table 2.

Priority	Code	Description in CyberAid context
Must Have	MH	Requirement is essential for the CyberAid baseline, architecture, regulatory alignment or pilot validation
Should Have	SH	Requirement is important and should be implemented if technically feasible within the planned work
Could Have	CH	Requirement adds value but is not critical for the first baseline
Won't Have	WH	Requirement is out of scope for the current version or postponed for later consideration

Table 2: MoSCoW Prioritisation Template

These priorities were then discussed with the relevant partners and could change as the architecture, components and pilot scenarios became more mature.

Where useful, requirements were also translated into user stories. User stories helped to express requirements in a more practical way, from the perspective of a user or actor interacting with the system. This approach was useful because it avoided describing everything only in technical terms and made it easier to connect requirements with validation scenarios. The user stories followed the template outlined in Table 3.

Field	Description
US-ID	Unique identifier for the user story
Title	Short descriptive title
Role(s)	Role or roles of involved actors
Business Value	Expected value or benefit provided by the user story
Dependencies & Assumptions	Dependencies with other user stories or requirements and Assumptions needed for the user story to be valid
Acceptance Criteria	What needs to be achieved for the user story to be considered fulfilled
Related Requirements	Requirement ID(s) from the requirements catalogue that are addressed, supported, or validated by the user story.

Table 3: User Stories Template

The user stories were expected to be mapped, where possible, to pilot scenarios and project tasks. This mapping would support later validation activities and would also help to show how the requirements were connected with the actual use of CyberAid in financial-sector environments.

2.4 Stakeholder Engagement and Validation Process

Stakeholder engagement formed part of the requirements methodology because the requirements needed to reflect both technical expectations and operational needs. For CyberAid, stakeholders included technical partners, pilot partners, partners working on regulatory and legal aspects, WP leaders and, where applicable, external financial-sector representatives involved in validation activities.

The first level of stakeholder engagement was performed through partner contributions to the shared requirements collection file. Each partner was expected to provide requirements linked to its tools, tasks, standards, legislation, pilot needs or cybersecurity expertise. This input was then reviewed and discussed with the relevant partners where clarifications were needed.

The second level of engagement took place through WP2 calls and requirements discussions. During these discussions, partners reviewed the collected requirements, clarified their meaning, identified missing inputs, discussed priorities and checked dependencies with architecture, CTI, monitoring, orchestration, reporting and integration work.

The third level of engagement concerned pilot and stakeholder validation. Pilot partners were expected to review the requirements and user stories linked to their operational scenarios. This was important in order to verify whether the proposed requirements were realistic, relevant and possible to validate during the project.

The validation of requirements was planned to be performed progressively. Requirements could be accepted, revised, merged, postponed or removed. Where possible, the reason for these changes would be recorded in order to maintain traceability between the original input, the final requirement, the related user story and the WP or pilot where it was relevant. The validation statuses used to track the maturity and review state of each requirement are presented in Table 4.

Status	Meaning
Draft	Requirement has been proposed but not yet reviewed
Under Review	Requirement is being checked
Accepted	Requirement is included in the baseline catalogue
To Be Revised	Requirement needs clarification or modification
Merged	Requirement has been merged with another similar requirement
Deferred	Requirement is valid but postponed for later project phases
Rejected	Requirement is not relevant or out of scope

Table 4: Requirement Validation Status

Figure 3 illustrates the iterative feedback loop followed during stakeholder validation. Candidate requirements are first drafted and reviewed, then discussed within WP2 and with relevant partners or pilot stakeholders. Based on the feedback received, each requirement may be updated and classified as accepted, revised or deferred.




Figure 3: Stakeholder Validation and Feedback Loop

The initial validation of the requirements baseline was linked with MS2.1, which concerns the completion of the initial cybersecurity requirements analysis for the financial sector. The validation evidence may include workshop minutes, approval documents, WP2 discussion outcomes, and partner feedback on the requirements catalogue.

3. FINANCIAL SECTOR CONTEXT AND THREAT LANDSCAPE

The financial sector is one of the most critical domains for the European economy and society, as it supports core services such as banking, payments, investments, trading, insurance, clearing and settlement. These services rely heavily on digital infrastructures, interconnected platforms, third-party service providers and real-time data exchanges. As a result, the cybersecurity posture of financial organisations is directly connected with the continuity and trustworthiness of the wider financial ecosystem.

In the context of CyberAid, the financial sector is considered not only as a group of individual financial organisations, but also as a set of critical infrastructures and digital services that depend on each other. A disruption in one part of this ecosystem may affect other organisations, customers, markets or public services. For this reason, financial cybersecurity requirements need to consider both the protection of individual systems and the resilience of broader financial processes.

Financial institutions are also operating in an environment where cyber threats are becoming more complex and more frequent. Threat actors target customer data, transaction systems, payment services, authentication mechanisms, APIs, cloud environments, internal networks and third-party connections. At the same time, the sector is under strong regulatory pressure, with frameworks such as DORA, NIS2, GDPR, PSD2, the EU AI Act and other financial-sector regulations introducing requirements related to operational resilience, incident reporting, data protection, ICT risk management and trusted use of AI.

The role of CyberAid in this context is to support the identification of cybersecurity requirements that are relevant for financial technologies and infrastructures, including critical infrastructures. The project focuses on improving cyber-resilience through an AI-driven framework that combines proactive defence, monitoring, incident response, reporting, privacy-preserving analytics, security knowledge modelling and intelligent orchestration.

This section provides the context for the requirements presented later in the deliverable. It explains the main characteristics of financial critical infrastructures, the digital transformation trends affecting the sector, and the cybersecurity challenges that need to be considered when defining the CyberAid requirements baseline.

3.1 Overview of the Financial Sector as Critical Infrastructure

Financial critical infrastructures include the systems, platforms, services and operational processes that allow financial organisations to provide essential services. These may include core banking systems, payment infrastructures, trading platforms, clearing and settlement systems, customer-facing digital channels, fraud detection systems, identity and access management services, data analytics platforms and connections with third-party providers.

The digital transformation of the financial sector has increased the dependency of financial organisations on software, data, cloud services, APIs and automated decision-making systems. Banks, payment institutions, asset management firms, FinTech organisations and other financial service providers are increasingly using digital platforms to deliver services faster and at larger scale. This creates benefits for efficiency and customer experience, but it also expands the attack surface that needs to be protected.

One important characteristic of financial infrastructures is their high level of interconnection. Financial organisations rarely operate in isolation. They exchange data with customers, regulators, payment networks, cloud providers, technology vendors, other financial institutions and external service providers. This means that cybersecurity requirements cannot only focus on internal systems, but also need to cover secure communication, third-party risk, API security, data exchange, supply chain dependencies and interoperability.

Another important aspect is the coexistence of legacy systems with modern technologies. Many financial organisations still rely on older infrastructures, while at the same time integrating cloud-native services,

AI-based analytics, open banking APIs, mobile banking platforms and automated monitoring tools. This mixed environment creates practical challenges for security monitoring, vulnerability management, integration, logging, incident response and compliance reporting.

AI and Generative AI are also becoming increasingly relevant in financial-sector operations and cybersecurity. These technologies can support fraud detection, anomaly detection, threat intelligence analysis, incident investigation and automated reporting. However, they also introduce new requirements related to explainability, human oversight, data protection, model security, robustness, accountability and prevention of misuse. For CyberAId, this is especially important because the project relies on AI-driven and LLM-based orchestration mechanisms to support cybersecurity operations.

Financial critical infrastructures therefore need cybersecurity capabilities that are not limited to perimeter protection. They require continuous monitoring, proactive vulnerability management, secure data handling, regulatory alignment, fast incident response and the ability to correlate information from multiple sources. These needs form the basis for the requirements analysis performed in this deliverable and are further examined in the following sections.

3.2 Financial-Sector Attack Surface and Key Attack Vectors

From a daily operations perspective, the financial industry has a broad and complex attack surface that is comprised of the number of systems, users, data flows, external dependencies and high-value transactions. Customer-facing applications, payment channels, APIs, trading platforms, identity and access management systems, internal networks, cloud services, third-party connections and operational monitoring tools are all targets for attackers.

The most relevant threat actors in the financial sector consist of cybercriminal groups, ransomware operators, fraud networks, state-sponsored or state-aligned actors, hacktivists, malicious insiders, and opportunistic attackers exploiting exposed services or misconfigurations. Their motivations may include financial gain, data theft, service disruption, espionage, fraud, extortion or reputational damage.

Common attack vectors include: phishing and social engineering, credential theft, account takeover, malware and ransomware, exploitation of vulnerable software, API abuse, supply chain compromise, distributed denial-of-service attacks, insider misuse, data leakage, cloud misconfiguration and manipulation of logs, transactions or monitoring data. In financial environments, these attack vectors can have a direct impact on the confidentiality, integrity, availability, regulatory compliance and customer trust.

From the CyberAId perspective, these attack vectors are relevant because they create requirements for proactive vulnerability management, continuous monitoring, anomaly detection, privacy-preserving analytics, threat intelligence correlation, incident response, regulatory reporting and human-in-the-loop decision support.

3.3 Emerging Threats in Finance

The financial sector remains one of the most exposed and systemically important domains in the contemporary cyber threat landscape. Financial entities process high-value transactions, manage sensitive personal and business data, support payment and settlement infrastructures, and provide services whose disruption can affect citizens, enterprises, markets, and wider critical infrastructure. Their technology environments are also highly complex, combining legacy banking systems, cloud services, payment infrastructures, trading platforms, customer-facing APIs, mobile and web applications, identity and access management services, data analytics pipelines, and interconnected third-party platforms. This digital expansion has improved service delivery and operational efficiency, but it has also increased the number of exploitable entry points available to malicious actors. ENISA's threat landscape analysis for the finance sector confirms that financial entities in Europe are affected by a broad range of cyber incidents, including

data breaches and leaks, fraud-related activity, disruption of services, supply chain compromise, and attacks against credit institutions [1].

The emerging threat landscape in finance is not defined only by the appearance of new attack techniques, but also by the convergence of existing threats with new technological, operational, and regulatory dependencies. Traditional attack vectors such as phishing, credential theft, malware, ransomware, and denial-of-service attacks remain highly relevant. However, these techniques are increasingly supported by automation, artificial intelligence, cloud exploitation, supply chain compromise, and more sophisticated fraud infrastructures. As a result, financial institutions are required to defend not only individual systems, but also end-to-end business processes, third-party dependencies, identity flows, data pipelines, transaction chains, and customer interaction channels. ENISA's 2024 threat landscape identifies threats against availability, ransomware, and threats against data among the most significant cybersecurity categories, all of which are directly relevant to financial services where continuity, integrity, and trust are core operational requirements [2].

A key driver of emerging risk is the growing digitalisation and interconnection of financial services. Open banking, embedded finance, mobile banking, digital onboarding, account aggregation, real-time payments, and fintech integrations depend on secure APIs, trusted identity mechanisms, reliable data exchange, and continuous service availability. These developments create new business opportunities, but they also introduce additional exposure to API abuse, account takeover, payment manipulation, weak authorisation controls, excessive data exposure, and fraud across interconnected platforms. In parallel, the increased reliance on cloud infrastructure and outsourced ICT services introduces new governance, configuration, identity, monitoring, and concentration risks. Financial organisations therefore need to manage cybersecurity not only as a perimeter defence problem, but as a continuous assurance challenge across distributed systems, external providers, and digital service chains.

The impact of disruptive cyberattacks is particularly severe in finance because availability is directly linked to economic activity and public confidence. Payment processing, online banking, trading operations, customer authentication, settlement processes, and regulatory reporting depend on uninterrupted service delivery and reliable performance. Ransomware, destructive malware, distributed denial-of-service attacks, and extortion campaigns may therefore evolve rapidly from technical incidents into business continuity crises. Modern ransomware operations are also increasingly multi-dimensional, combining encryption, data theft, operational disruption, reputational pressure, and threats of public disclosure. For financial entities, this creates simultaneous risks to confidentiality, availability, regulatory compliance, customer trust, and operational resilience.

Beyond service disruption and data theft, financial institutions must also address threats to data integrity. In several financial scenarios, the most damaging attack may not be the exfiltration of data, but its subtle manipulation. Attackers may attempt to alter transaction records, modify beneficiary details, manipulate risk indicators, tamper with trading signals, poison fraud detection datasets, or corrupt logs used for investigation and compliance. These attacks may be difficult to detect because affected systems can remain operational while producing incorrect, misleading, or incomplete outputs. In finance, the consequences can include fraudulent payments, inaccurate risk assessment, market abuse, failed reconciliation, weakened auditability, and loss of evidentiary value during incident response.

The regulatory environment further increases the importance of understanding emerging threats in a structured way. DORA, NIS2, GDPR, PSD2, and the EU AI Act reinforce the need for stronger governance, accountability, resilience testing, third-party risk management, data protection, and transparency in automated systems. DORA is particularly relevant because it consolidates ICT risk management, ICT-related incident reporting, digital operational resilience testing, and ICT third-party risk management into a common framework for the financial sector [3]. This means that emerging threats must be translated into concrete cybersecurity requirements, not only for prevention and detection, but also for incident classification, evidence collection, regulatory reporting, and resilience validation.

Within this context, three emerging threat domains require particular attention in the financial sector. First, Artificial Intelligence and Generative AI are changing the scale, realism, and automation level of cyberattacks and fraud, while also introducing new vulnerabilities in AI-enabled financial systems. Second, supply chain attacks are becoming more critical as financial entities increasingly depend on cloud platforms, ICT providers, fintech integrations, managed services, and open-source software components. Third, post-quantum cryptography introduces a strategic long-term resilience challenge, as current cryptographic mechanisms protecting financial transactions, communications, authentication processes, digital signatures, and sensitive records may become vulnerable to future quantum capabilities. These three areas are analysed in the following subsections as key drivers for financial cybersecurity requirements.

3.3.1 GDPR AI/GenAI-related Vulnerabilities and Threats

Artificial Intelligence (AI) and Generative AI (GenAI) are becoming increasingly relevant to the cybersecurity posture of financial institutions. Their adoption supports fraud detection, transaction monitoring, anti-money laundering analysis, credit risk assessment, customer support, behavioural analytics, security operations, and automated decision support. At the same time, these technologies introduce a new class of vulnerabilities and amplify existing cyber and fraud threats. In the financial sector, this dual role is particularly important: AI can strengthen detection and response capabilities, but it can also expand the attacker's toolkit, increase the scale of social engineering, and create new weaknesses in AI-enabled systems. NIST's Generative AI Profile under the AI Risk Management Framework highlights that GenAI risks must be managed across the full AI lifecycle, considering design, development, deployment, monitoring, misuse, and organisational governance [4].

One of the most immediate risks is the use of GenAI to improve social engineering and fraud operations. Financial institutions are already exposed to phishing, business email compromise, investment fraud, account takeover, authorised push payment fraud, and impersonation attacks. GenAI increases the credibility, speed, and personalisation of these campaigns. Attackers can generate fluent multilingual messages, imitate the writing style of executives or customers, create realistic fake support conversations, and adapt fraud content to specific institutions, geographies, or customer profiles. Europol's IOCTA 2024 notes that AI tools expand the possibilities for fraudsters by enabling them to reach more victims and improve social engineering techniques [5]. For financial organisations, this affects customer communication, call-centre interactions, payment authorisation workflows, fraud investigation teams, and internal approval processes.

Deepfake-enabled impersonation further increases this exposure. AI-generated voice, video, and synthetic identity artefacts can be used to manipulate employees, bypass weak verification steps, or support fraudulent onboarding. The risk is not limited to external customer fraud. It may also affect internal operations, especially where payment approvals, changes to beneficiary accounts, urgent fund transfers, or sensitive information requests depend on human judgement and trust-based communication. A realistic voice clone of a senior executive, a fabricated video call, or synthetic documentation may be sufficient to trigger a fraudulent process when controls rely only on visual or verbal confirmation. This creates a hybrid threat that combines cybercrime, identity fraud, business process compromise, and psychological manipulation.

AI also changes the technical capabilities of attackers. GenAI systems can support reconnaissance, vulnerability discovery, malware development, exploit adaptation, and automation of attack chains. Although current models may not independently conduct sophisticated attacks without human guidance, they can reduce the skill barrier for less experienced actors and accelerate parts of the cybercrime workflow. This is relevant for financial institutions because attackers can use AI to scan public-facing applications, identify weak APIs, draft phishing infrastructure, generate scripts, summarise leaked data, or automate credential-stuffing support tasks. ENISA's finance-sector threat landscape already identifies fraud, data breaches, service disruption, and supply chain compromise as relevant threats for financial entities, and AI can act as an accelerator across these categories [1].

In parallel, financial institutions that deploy AI-enabled systems introduce their own AI-specific vulnerabilities. These include adversarial examples, data poisoning, model poisoning, model inversion, membership inference, model extraction, prompt injection, insecure output handling, excessive agency, and leakage of sensitive information. In transaction monitoring or fraud detection, adversarial manipulation may allow attackers to craft behaviours that evade detection. In AI-assisted customer service, prompt injection may cause a chatbot or assistant to ignore instructions, reveal sensitive information, or interact unsafely with connected tools. In AI-supported security operations, poisoned data may distort detection logic, generate misleading alerts, or reduce analyst trust in automated outputs. OWASP's Top 10 for Large Language Model Applications identifies prompt injection, sensitive information disclosure, supply chain vulnerabilities, data and model poisoning, improper output handling, excessive agency, and system prompt leakage among the key risks affecting LLM-based applications [6].

The use of AI in finance also creates risks linked to data governance and confidentiality. AI models may process personal data, transaction histories, behavioural profiles, financial records, customer complaints, security logs, and incident information. If such data is used without proper controls, there is a risk of privacy violations, unauthorised disclosure, retention of sensitive information, or leakage through model outputs. This is particularly significant where external AI services or third-party model providers are used. Financial entities must ensure that training, fine-tuning, retrieval-augmented generation, logging, and inference processes do not expose confidential business data, customer data, payment information, credentials, or security-sensitive artefacts. These risks are not only technical, but also regulatory, given the obligations imposed by GDPR, DORA, and the EU AI Act.

Another important concern is the reliability and explainability of AI-generated outputs. In financial cybersecurity operations, AI may be used to support alert prioritisation, incident triage, fraud scoring, threat intelligence summarisation, or recommended response actions. If the system produces inaccurate, incomplete, biased, or hallucinated outputs, analysts may take incorrect decisions or overlook important evidence. This risk becomes more serious when AI systems are connected to operational tools, ticketing systems, identity platforms, or automated response mechanisms. Improper output handling can also create security vulnerabilities when generated content is passed directly into downstream systems without validation or sanitisation [6]. Therefore, human oversight, auditability, confidence scoring, validation, and controlled levels of automation remain essential.

AI-related threats are also connected to supply chain and dependency risks. Financial institutions may rely on pre-trained models, external AI APIs, open-source libraries, vector databases, plugins, datasets, and cloud-hosted AI services. Each component introduces potential exposure to compromise, manipulation, dependency vulnerabilities, licensing constraints, or insufficient transparency. A compromised model, poisoned dataset, insecure plugin, or vulnerable AI library may affect the integrity of the entire AI-enabled workflow. This is particularly relevant in finance, where AI components may influence fraud detection, security monitoring, customer interaction, or compliance-related processes.

For financial entities, the main challenge is that AI-related risk cannot be treated as a single technical issue. It affects fraud prevention, identity verification, cybersecurity operations, data protection, third-party governance, and business process integrity at the same time. This means that controls around AI systems must be embedded before deployment and continuously reassessed during operation. Secure AI use in finance requires clear ownership of AI-enabled processes, validation of training and operational data, monitoring for adversarial or abnormal behaviour, protection of sensitive information used by models, and strict control over the actions that AI systems are allowed to recommend or execute. These measures are particularly important where AI outputs influence security alerts, transaction monitoring, customer interaction, or incident response decisions.

3.3.2 Supply Chain Attacks

Supply chain attacks have become a central concern for the financial sector because financial entities rarely operate as isolated organisations. Banks, payment institutions, investment firms, insurance providers,

trading platforms, fintech companies, and market infrastructure operators rely on a broad ecosystem of ICT providers, cloud platforms, managed service providers, software vendors, identity providers, data processors, consultants, open-source libraries, and outsourced operational services. This dependency enables financial organisations to scale, innovate, and integrate new digital services, but it also means that risks may enter the organisation through trusted third parties rather than through direct attacks against internal systems. ENISA's finance-sector threat landscape explicitly identifies supply chain attacks among the methods used by threat actors to exploit vulnerabilities for financial gain, alongside fraud and social engineering [1].

In the financial sector, the importance of supply chain security is amplified by the nature of the services being protected. A compromised supplier may affect customer authentication, payment processing, transaction monitoring, regulatory reporting, cloud infrastructure, trading connectivity, or security monitoring. This makes supply chain risk both a cybersecurity issue and an operational resilience issue. The impact is not limited to confidentiality or data theft; it may also include service unavailability, manipulation of transaction flows, loss of auditability, unauthorised access to customer data, or disruption of critical financial processes. As financial services become more interconnected, attacks against one provider may have cascading consequences across several institutions that depend on the same technology or service.

A common form of supply chain attack involves the compromise of software development and delivery processes. Financial institutions depend on commercial software, internally developed applications, open-source components, APIs, CI/CD pipelines, container images, software updates, and third-party integrations. Attackers may attempt to insert malicious code into software updates, compromise build environments, abuse developer credentials, poison open-source dependencies, or exploit vulnerabilities in widely used libraries. Such attacks are difficult to detect because malicious components may be delivered through trusted channels and may appear as legitimate updates or dependencies. CISA highlights software transparency, including the use of Software Bills of Materials, as a key mechanism for understanding software composition and managing known risks across the software supply chain [7].

Open-source software introduces a particular challenge. Financial entities often rely on open-source libraries and frameworks for application development, data processing, cryptography, monitoring, DevOps automation, and AI-enabled services. Although open-source software can support transparency and innovation, it may also introduce risks when components are unmaintained, poorly reviewed, vulnerable, or replaced by malicious packages. Dependency confusion, typosquatting, malicious package publication, and compromise of maintainers are realistic attack paths. These risks are especially important in financial environments where even minor software components may become part of systems supporting payments, customer data, fraud detection, or regulatory workflows.

Cloud and managed service dependencies represent another major dimension of supply chain risk. Financial institutions increasingly use cloud platforms for infrastructure, analytics, storage, collaboration, cybersecurity tooling, and customer-facing services. They may also outsource network operations, incident response support, identity management, vulnerability scanning, helpdesk functions, or compliance tooling. These dependencies can improve efficiency and resilience when properly governed, but they also create concentration risk and shared failure modes. If a critical cloud provider, managed service provider, or identity provider suffers a breach or outage, the impact may extend beyond a single financial institution. DORA directly addresses this concern by establishing an EU-wide oversight framework for critical ICT third-party providers and by requiring financial entities to manage ICT third-party risk as part of their ICT risk management framework [8], [3].

The risk is not only technical but also contractual and governance-related. Financial entities may lack sufficient visibility into subcontractors, fourth parties, data processing locations, incident notification obligations, audit rights, vulnerability management practices, or security controls applied by suppliers. This lack of transparency can delay detection, limit incident response, and complicate regulatory reporting. The European Central Bank's guidance on outsourcing cloud services stresses that supervised entities should apply an appropriate level of diligence to outsourced ICT services, including security, risk management, and

control expectations comparable to those applied to internal services [9]. This is particularly relevant where outsourced services support critical or important functions.

Supply chain attacks may also target business and operational processes rather than software alone. Examples include compromise of a payment processor, abuse of a fintech API integration, manipulation of a third-party data feed, compromise of a customer communication provider, or exploitation of a service provider with privileged remote access. In these cases, the attacker may use the supplier relationship as a trusted pathway into financial operations. The threat may be difficult to identify because activity originating from a supplier may be treated as legitimate by default. Strong authentication, least-privilege access, segmentation, logging, behavioural monitoring, and continuous validation of third-party access are therefore essential controls.

The regulatory relevance of supply chain security has increased significantly in Europe. DORA requires financial entities to manage ICT third-party risk, maintain appropriate contractual arrangements, monitor ICT providers, and address risks linked to outsourcing and concentration. The Bank of Greece summarises DORA's main requirement areas as ICT risk management, ICT-related incident management, digital operational resilience testing, and management of ICT third-party risk [10]. This confirms that supply chain attacks are not only a technical threat category, but also a core requirement driver for operational resilience in the financial sector.

For financial institutions, supply chain security must therefore be approached as a continuous assurance process. It requires maintaining visibility over critical suppliers and software components, assessing supplier security before and during engagement, monitoring third-party access, validating software integrity, managing vulnerabilities in dependencies, and ensuring that incident response processes include supplier coordination. In practice, the ability to identify which financial services depend on which ICT providers, software components, APIs, cloud services, and subcontractors becomes essential for both cyber defence and regulatory resilience.

3.3.3 Post Quantum Cryptography

Post-Quantum Cryptography (PQC) is becoming an important strategic consideration for the financial sector because modern financial services depend extensively on cryptographic mechanisms that may be affected by future quantum computing capabilities. Public-key cryptography is used across online banking, payment authorisation, interbank communication, secure APIs, customer authentication, digital signatures, certificates, secure software updates, encrypted communications, and long-term protection of sensitive financial and personal data. Although cryptographically relevant quantum computers are not yet available at operational scale, their expected ability to weaken or break widely used asymmetric cryptographic algorithms creates a long-term risk that financial institutions must begin to address in advance. The Bank for International Settlements (BIS) stresses that the financial system should start preparing now for the transition to quantum-safe cryptographic infrastructures, with awareness and cryptographic inventory as critical foundations [11].

The core concern is that quantum computing may undermine the security assumptions of currently deployed public-key algorithms. Many financial systems rely on RSA, Diffie-Hellman, and Elliptic Curve Cryptography for key exchange, authentication, digital signatures, and secure communications. These algorithms are considered secure today because classical computers cannot efficiently solve the mathematical problems on which they are based. However, sufficiently powerful quantum computers running algorithms such as Shor's algorithm could make several of these mechanisms vulnerable. This would affect not only encrypted communications, but also the authenticity and integrity guarantees provided by digital signatures and certificates. For financial entities, the consequences may include exposure of confidential data, weakened customer authentication, compromised transaction integrity, invalidated digital signatures, and reduced trust in secure communication channels.

A particularly relevant threat is the "harvest now, decrypt later" scenario. In this model, adversaries collect encrypted data today, even if they cannot currently decrypt it, with the expectation that future quantum

capabilities will allow them to recover the plaintext. This is a serious concern for financial institutions because many types of data retain value for long periods. Examples include customer identity records, payment histories, trading data, contracts, legal documents, business intelligence, authentication records, and regulatory evidence. Where data confidentiality must be preserved for several years or decades, quantum risk becomes relevant before quantum computers are actually available. NIST's guidance on the transition to post-quantum cryptography recognises that some use cases may require earlier migration to quantum-resistant key establishment schemes because of this long confidentiality lifetime [12].

The financial sector also faces a challenge because cryptographic mechanisms are deeply embedded in complex and distributed infrastructures. Cryptography is not limited to visible systems such as internet banking portals or payment gateways. It is also present in hardware security modules, smart cards, mobile applications, ATM networks, interbank messaging, internal APIs, databases, backup systems, code signing, software update mechanisms, VPNs, TLS certificates, identity federation, logging systems, and machine-to-machine communication. In many organisations, the exact location and purpose of cryptographic assets are not fully documented. This lack of visibility makes migration planning difficult. A financial entity cannot move toward quantum-safe cryptography effectively unless it first understands which algorithms, protocols, keys, certificates, libraries, and systems are currently in use.

This is why cryptographic inventory and crypto-agility are central to post-quantum readiness. A cryptographic inventory provides structured knowledge of where cryptography is used, which algorithms are deployed, what data or processes they protect, and when associated keys or certificates expire. Crypto-agility refers to the ability to replace or update cryptographic algorithms, protocols, and implementations without major redesign of business-critical systems. In financial environments, this capability is essential because migration will not occur at once. Different systems will have different risk levels, vendor dependencies, certification requirements, operational constraints, and performance impacts. Legacy systems, payment infrastructures, trading platforms, and regulated reporting systems may require careful testing before cryptographic changes can be deployed.

The standardisation landscape has advanced significantly. In August 2024, NIST released the first three finalised post-quantum cryptography standards: FIPS 203 for ML-KEM, intended for general encryption and key establishment; FIPS 204 for ML-DSA, intended for digital signatures; and FIPS 205 for SLH-DSA, another digital signature standard based on a different mathematical approach [13]. These standards provide an important technical foundation for future migration. At the same time, ENISA has highlighted that PQC integration is not only a matter of selecting algorithms, but also of integrating them into existing protocols, applications, infrastructures, and operational environments [14]. This is particularly important for finance, where systems are highly regulated, performance-sensitive, and dependent on interoperability across institutions and providers.

In Europe, the policy direction also points toward coordinated preparation. The European Commission Recommendation on a coordinated implementation roadmap for the transition to post-quantum cryptography encourages Member States to develop comprehensive strategies for PQC adoption and to support a synchronised transition across public administrations and critical infrastructures [15]. A coordinated approach is especially relevant for the financial sector because financial services depend on cross-border communication, shared infrastructures, supervisory expectations, and common trust frameworks. If institutions migrate at different speeds or adopt incompatible approaches, interoperability and operational stability may be affected.

PQC migration also creates operational and governance challenges. New algorithms may have different key sizes, signature sizes, computational requirements, bandwidth implications, and implementation risks compared with existing mechanisms. These characteristics may affect latency-sensitive financial systems, mobile banking applications, embedded devices, smart cards, hardware security modules, and high-frequency trading environments. Migration must therefore be supported by testing, performance assessment, vendor coordination, rollback planning, and risk-based prioritisation. Hybrid approaches, where classical and post-quantum algorithms are temporarily combined, may be useful during transition phases, but they also increase implementation complexity and must be carefully managed.

For financial entities, post-quantum cryptography should therefore be treated as a long-term cyber resilience and risk management issue rather than as a purely technical cryptographic upgrade. The immediate priorities are to identify cryptographic dependencies, assess which data and services require long-term protection, monitor standards and regulatory expectations, engage critical ICT providers, and define migration roadmaps for high-risk systems. This is particularly important for services involving customer authentication, payment authorisation, secure interbank communication, digital signatures, certificate management, software integrity, and confidential data exchange. Early preparation reduces the risk of rushed, fragmented, or incompatible migration when quantum-safe transition becomes operationally mandatory.

3.4 Threat Intelligence Landscape

Cyber Threat Intelligence (CTI) is a fundamental input for understanding the cybersecurity requirements of the financial sector. Financial entities operate in a threat environment characterised by financially motivated cybercrime, ransomware, fraud, geopolitical activity, hacktivism, data breaches, supply chain compromise, and attacks against critical digital infrastructure. Banks, payment institutions, investment firms, fintech providers, insurance organisations, trading platforms, and financial market infrastructures must therefore maintain continuous awareness of the actors, tactics, techniques, vulnerabilities, campaigns, and indicators that may affect their operations. In this context, CTI supports not only security monitoring, but also threat modelling, risk assessment, incident response planning, third-party risk analysis, resilience testing, and regulatory preparedness.

The value of CTI in the financial sector depends on the ability to combine several levels of intelligence. Strategic intelligence provides an understanding of sectoral trends, threat actor motivations, geopolitical developments, and systemic risks. Tactical intelligence focuses on adversary tactics, techniques, and procedures, allowing security teams to understand how attacks are executed. Operational intelligence supports active investigations, campaigns, and incident response, while technical intelligence provides concrete artefacts such as malicious IP addresses, domains, file hashes, malware signatures, phishing infrastructure, exploit references, and vulnerability information. For financial institutions, these layers should not be treated separately. A technical indicator becomes more useful when it can be connected to a known campaign, a relevant business process, an affected asset, a vulnerability, or a specific operational impact.

The financial-sector CTI ecosystem includes several sources of intelligence. These include public and sector-specific threat reports, national and European cybersecurity agencies, financial information sharing communities, Computer Security Incident Response Teams (CSIRTs), vulnerability databases, malware repositories, fraud intelligence sources, open-source intelligence, dark web monitoring, commercial intelligence feeds, and internal telemetry from SIEM, XDR, endpoint, network, identity, and transaction-monitoring systems. ENISA's finance-sector threat landscape provides a European-level view of the threat environment affecting financial entities, including data breaches, service disruption, fraud, supply chain attacks, and attacks against credit institutions [1]. FS-ISAC also provides financial-sector analysis and information sharing, highlighting risks such as ransomware, DDoS, fraud, third-party compromise, geopolitical activity, and the need for trusted exchange of intelligence across the sector [16].

A key requirement for effective CTI is standardisation. Financial institutions often receive intelligence from multiple sources and in different formats, which can create fragmentation and reduce operational value. Structured standards such as STIX and TAXII support the consistent exchange of CTI by representing threat actors, malware, campaigns, attack patterns, indicators, vulnerabilities, tools, and observed data in a common format [17], [18]. At the level of D2.1, these standards are relevant because they demonstrate the need for interoperable and reusable threat information. Their detailed application within CyberAid's CTI integration mechanisms and knowledge modelling activities is addressed in D2.4.

MITRE ATT&CK is another important reference for the financial-sector threat intelligence landscape. ATT&CK provides a globally accessible knowledge base of adversary tactics and techniques based on real-world observations [19]. For financial entities, its value lies in translating threat intelligence from isolated indicators into behavioural patterns. Instead of only blocking known malicious infrastructure, security teams can map adversary activity to tactics such as initial access, execution, persistence, privilege escalation, defence evasion, credential access, discovery, lateral movement, collection, command and control, exfiltration, and impact. This supports detection engineering, control validation, red-team planning, security testing, incident triage, and gap analysis.

MITRE has also developed an enhanced cyber threat model for financial services sector institutions, reflecting attacker methods relevant to financial organisations [20]. This is particularly useful because attacks against financial entities often combine technical intrusion with fraud, identity abuse, process manipulation, data exfiltration, and disruption of financial services. For example, credential abuse may support account takeover or unauthorised access to payment workflows. API reconnaissance may precede data collection or transaction manipulation. Malware activity affecting payment systems may involve persistence, command and control, and impact-oriented techniques. Ransomware may combine initial access, privilege escalation, lateral movement, data theft, and service disruption. ATT&CK-based mapping therefore helps financial institutions understand how individual events may form part of a broader adversary sequence.

Threat intelligence also has an important role in deriving cybersecurity requirements. Intelligence about threat actors, attack techniques, exploited vulnerabilities, targeted financial services, and observed incident patterns can inform requirements for monitoring, detection, access control, logging, resilience testing, incident response, fraud prevention, and third-party assurance. For example, intelligence on credential theft and account takeover can inform stronger identity monitoring and behavioural analytics requirements. Intelligence on ransomware can inform backup, segmentation, recovery, and incident escalation requirements. Intelligence on supply chain attacks can inform supplier assurance, software integrity, SBOM, and third-party monitoring requirements. In this way, CTI acts as a bridge between the external threat environment and the internal security controls required by financial entities.

At the same time, financial CTI must preserve sufficient context to remain actionable. The same vulnerability, indicator, or attack technique may have different relevance depending on whether it affects an internal reporting tool, a payment gateway, a trading interface, an identity provider, or a customer-facing banking application. For this reason, CTI should be interpreted in relation to asset criticality, business process dependencies, exposure level, regulatory relevance, and incident history. This contextual interpretation is especially important for prioritisation, because financial institutions receive large volumes of intelligence that cannot all be treated with the same urgency.

In the context of CyberAId, the threat intelligence landscape provides input for subsequent knowledge-modelling and orchestration activities. Structured CTI, including ATT&CK mappings, STIX/TAXII artefacts, vulnerability information, and sector-specific threat reports, can support the later representation of threats, assets, controls, incidents, and operational evidence in dedicated security knowledge models. The detailed specification of these models, their semantic structure, and their CTI integration mechanisms is addressed in D2.4, while D2.1 uses CTI primarily as a source for understanding financial-sector cybersecurity requirements.

4. REGULATORY AND STANDARDS FRAMEWORK

The development and deployment of CyberAid takes place within a dense and evolving regulatory environment that governs cybersecurity data protection operational resilience and the use of artificial intelligence across the European financial sector. Understanding this environment is not merely a compliance exercise. It is a foundational requirement for designing cybersecurity solutions that are credible adoptable and sustainable for financial institutions operating under direct supervisory scrutiny.

This section provides a structured overview of the regulatory instruments and technical standards that directly shape the requirements architecture and validation criteria of the CyberAid platform. The frameworks described here constitute the compliance baseline against which the project's deliverables are assessed and against which the pilot deployments will be evaluated. A more detailed legal analysis of each instrument including article-level obligation mapping and gap assessment is provided in Deliverable D1.5.

4.1 EU Regulatory Framework Overview

The European regulatory landscape for financial cybersecurity is characterised by a layered architecture of directly applicable regulations and implemented directives each addressing distinct but interrelated aspects of digital resilience data protection payment security and technology governance. This layering creates a compliance environment of considerable complexity for financial institutions which must simultaneously satisfy obligations arising from instruments designed with different regulatory objectives supervisory communities and enforcement timelines in mind.

The primary instruments of relevance to CyberAid are the Digital Operational Resilience Act the NIS2 Directive the General Data Protection Regulation the revised Payment Services Directive and its associated technical standards the Fourth Anti-Money Laundering Directive and the EU AI Act. While each instrument addresses distinct regulatory concerns they interact extensively in practice and compliance with one frequently requires coordination with obligations arising under the others. CyberAid is designed to support financial institutions in navigating this complexity providing capabilities that address the overlapping requirements of these instruments through a coherent and integrated platform.

4.1.1 DORA (Digital Operational Resilience Act)

The Digital Operational Resilience Act (DORA), formally referred to as the EU Digital Operational Resilience Act, is a cornerstone of the European Union's strategy to strengthen the digital resilience of the financial sector. Established under Regulation (EU) 2022/2554[8], it introduces a unified regulatory framework designed to ensure that financial entities can effectively manage, withstand, and recover from information and communication technology (ICT) disruptions and cyber threats. DORA applies across banks, insurance companies, investment firms, and critical ICT third-party service providers, reflecting the increasing interdependence between financial services and digital infrastructure.

DORA was introduced in response to the growing digitalisation of financial services and the increasing dependence on complex ICT systems and external technology providers. Prior to its adoption, regulatory approaches to ICT risk varied significantly across EU Member States, resulting in fragmentation and inconsistent levels of operational resilience. The regulation addresses this gap by harmonising requirements and establishing a common supervisory framework across the EU financial ecosystem.

At its core, DORA addresses systemic risks arising from digital transformation, cloud adoption, and the concentration of critical services among a limited number of ICT providers. As highlighted by supervisory authorities such as the European Insurance and Occupational Pensions Authority (EIOPA)[3][21], the regulation is designed to strengthen operational resilience across the entire financial ecosystem by ensuring that ICT risk is managed not only at the level of individual institutions but also across interconnected supply chains.

A central pillar of DORA is the ICT risk management framework. Financial entities are required to implement robust governance structures, continuous risk identification processes, and comprehensive

internal controls covering all ICT assets and dependencies. This includes the ongoing monitoring of systems, applications, and external services to ensure that risks are identified, assessed, and mitigated in a timely manner.

DORA also introduces strict requirements for ICT-related incident management and reporting. Financial entities must detect, classify, and report major ICT incidents to competent authorities within defined timelines. This framework enhances regulatory transparency, enables cross-border coordination, and supports the monitoring of systemic cyber risks within the EU financial sector.

Another key pillar of DORA is digital operational resilience testing. The regulation mandates regular testing of ICT systems, including advanced threat-led penetration testing for critical functions. These exercises are designed to evaluate the ability of organisations to identify vulnerabilities and respond effectively to sophisticated cyberattacks, ensuring that resilience capabilities are continuously validated under realistic conditions.

In addition, DORA places strong emphasis on third-party ICT risk management. Financial institutions increasingly rely on external providers such as cloud platforms, data processors, and infrastructure vendors, which introduces additional systemic dependencies. The regulation therefore requires institutions to assess, monitor, and manage third-party risks, while also establishing an EU-level oversight framework for critical ICT service providers to reduce concentration risk across the financial system.

DORA further promotes information sharing among financial entities, encouraging the exchange of cyber threat intelligence and best practices. This collective approach aims to strengthen sector-wide awareness and improve the ability of financial institutions to respond to emerging threats in a coordinated manner.

Recent supervisory guidance, including materials published by European Banking Authority and BaFin on ICT outsourcing, cloud dependencies, and cyber resilience testing[22][36][37], further reinforces the importance of contractual safeguards, audit rights, and continuous monitoring of outsourced ICT services as part of DORA-aligned compliance practices.

Overall, DORA represents a shift from fragmented and reactive ICT risk management approaches towards a harmonised, proactive, and resilience-oriented regulatory model. By integrating risk management, incident reporting, resilience testing, third-party oversight, and information sharing into a single framework. As a result, it significantly strengthens the ability of the EU financial sector to maintain operational continuity and stability in the face of increasingly sophisticated cyber threats and growing digital interdependencies.

4.1.2 NIS2 Directive (Network and Information Security Directive 2)

The NIS2 Directive established under Directive (EU) 2022/2555[23] and replacing the original NIS Directive of 2016 constitutes the primary horizontal cybersecurity instrument applicable to operators of essential and important services across the European Union. For financial entities that fall within the scope of DORA the relationship between the two instruments is governed by a *lex specialis* arrangement: DORA takes precedence for ICT risk management and incident reporting obligations but NIS2 continues to apply in areas not covered or not addressed with equivalent specificity by DORA including supply chain security requirements management body accountability and certain aspects of vulnerability disclosure.

NIS2 significantly expands the scope of the original directive extending mandatory cybersecurity obligations to a larger number of sectors and to medium-sized entities that were previously excluded. It introduces a two-tier classification of essential and important entities with differentiated supervisory regimes and enforcement powers applicable to each category. Financial sector entities including credit institutions market infrastructure operators and payment service providers are classified as essential entities and therefore subject to the more stringent requirements of the upper tier.

The cybersecurity risk management measures mandated under Article 21 of NIS2 require entities to adopt technical and organisational measures proportionate to the risks they face covering a defined set of control

domains including risk analysis and information system security policies incident handling business continuity and crisis management supply chain security network security and access control and the use of multi-factor authentication. These requirements translate directly into security architecture specifications for CyberAid-BLUEPRINTS which provides reference implementations tailored to the financial sector context.

Supply chain security deserves particular attention in the NIS2 context. Article 21(2)(d) requires entities to address cybersecurity risks in their relationships with direct suppliers and service providers taking into account the overall quality of products and cybersecurity practices of those suppliers. This requirement read alongside DORA's third-party risk management framework creates a comprehensive compliance obligation that extends well beyond an institution's own perimeter. CyberAid addresses this through automated validation of API integrations with fintech partners and continuous monitoring of third-party access behaviour patterns.

NIS2 also introduces meaningful management body accountability provisions under Article 20. Management bodies of essential entities are required to approve cybersecurity risk management measures oversee their implementation and can be held personally liable for infringements resulting from a failure to comply. This governance dimension reinforces the importance of audit trails documented decision-making processes and human oversight mechanisms all of which are core design principles of the CyberAid platform.

4.1.3 GDPR (General Data Protection Regulation)

The General Data Protection Regulation established under Regulation (EU) 2016/679 [24] applies horizontally across all CyberAid data processing activities that involve personal data. Security monitoring anomaly detection insider threat identification and forensic investigation all involve to varying degrees the processing of personal data including authentication logs communication metadata user behavioural patterns and transaction records. GDPR therefore operates as a cross-cutting compliance constraint that must be integrated into the design of every CyberAid component that handles such data.

The core GDPR principles of purpose limitation and data minimisation require that personal data collected for security purposes is not repurposed for other uses and that only the data strictly necessary for the defined security objective is collected and retained. CyberAid-MONITOR is designed with configurable retention policies that align with both security requirements and regulatory data minimisation obligations and processing activities are scoped to clearly defined security purposes documented in the project's data management plan.

A particularly significant challenge in the financial cybersecurity context arises from the need to share threat intelligence across institutional boundaries without exposing sensitive customer or operational data. CyberAid addresses this through the privacy-preserving federated analytics framework implemented via PySyft [25] which enables collaborative fraud detection model training across financial institutions using federated learning and secure multi-party computation. Under this architecture raw data never leaves the originating institution and model updates are subject to differential privacy mechanisms before aggregation providing mathematical privacy guarantees that support GDPR compliance while enabling collective intelligence.

Data Protection Impact Assessments under Article 35 GDPR are required for processing activities that are likely to result in high risk to individuals including systematic monitoring of publicly accessible areas and large-scale processing of sensitive data. The deployment of security monitoring tools within financial institutions qualifies for DPIA in most cases and CyberAid's pilots will be accompanied by DPIA documentation developed in collaboration with the project's legal and ethics partner Fairlex.

Where security incidents processed through CyberAid involve personal data breaches the 72-hour supervisory notification obligation under Article 33 GDPR is triggered in parallel with applicable DORA incident reporting requirements. CyberAid-REPORT supports coordinated compliance with both

instruments through automated evidence collection scope and impact assessment and report generation capabilities reducing the operational burden on security teams managing dual regulatory timelines.

4.1.4 PSD2 and Strong Customer Authentication

The revised Payment Services Directive established under Directive (EU) 2015/2366 [26] and supplemented by Commission Delegated Regulation (EU) 2018/389 [27] on Strong Customer Authentication and secure open standards of communication establishes security requirements specifically applicable to payment service providers operating within the European Economic Area. PSD2 represents a foundational instrument for two of CyberAid's pilot deployments particularly the Anti-Money Laundering detection pilot with OTE/COSMOTE Payments and the banking cybersecurity pilot with CXB.

At the operational level PSD2 creates extensive security obligations around payment processing including requirements for transaction risk analysis fraud rate monitoring and the maintenance of Strong Customer Authentication for payment initiation. Payment service providers are required to submit annual assessments of operational and security risks to their competent authority under Article 96 PSD2 and these assessments must demonstrate the adequacy of the institution's risk mitigation measures. CyberAid-REPORT drawing on outputs from the OLISTIC risk assessment engine supports the preparation of these assessments by providing continuously updated risk scoring and mitigation evidence.

The open banking framework introduced by PSD2 creates extensive API ecosystems connecting payment service providers with third-party providers authorised to access customer account information and initiate payments. These APIs represent a significant and growing attack surface as they expose core banking functions to a wide range of external parties operating under varying security standards. CyberAid-PROACTIVE addresses this through continuous API traffic monitoring OAuth token validation and detection of abnormal API usage patterns that may indicate credential theft account takeover or data exfiltration attempts targeting open banking infrastructure.

The transaction monitoring capabilities embedded in CyberAid-MONITOR also interact directly with PSD2's fraud rate thresholds. Payment service providers can benefit from SCA exemptions for low-risk transactions only where they can demonstrate that their fraud rates remain below the prescribed thresholds under the Regulatory Technical Standards. Maintaining eligibility for these exemptions requires continuous monitoring of payment fraud patterns which is precisely the function served by CyberAid-MONITOR's machine learning-based anomaly detection for payment flows.

4.1.5 4AMLD (4th Anti-Money Laundering Directive)

The Fourth Anti-Money Laundering Directive established under Directive (EU) 2015/849 [28] and amended by the Fifth Anti-Money Laundering Directive (Directive (EU) 2018/843) establishes the core European framework for the prevention of money laundering and terrorist financing. For financial entities subject to both DORA and 4AMLD the two regulatory domains increasingly converge in practice: sophisticated money laundering operations routinely exploit cybersecurity vulnerabilities and the transaction data generated by AML monitoring systems provides a rich source of intelligence for broader threat detection purposes.

The central obligation imposed by 4AMLD on obliged entities is the requirement to monitor customer transactions on a continuous basis and to report suspicious transactions to the relevant Financial Intelligence Unit in their Member State. This monitoring obligation requires financial institutions to maintain detailed customer due diligence records to establish behavioural baselines for each customer relationship and identify deviations that may indicate money laundering activity. The practical challenge is substantial: traditional rule-based AML systems generate very high volumes of false positive alerts often exceeding 95 percent of total alerts in payment service environments while still missing sophisticated laundering techniques that are deliberately structured to remain below detection thresholds.

CyberAid addresses this challenge through a multi-layered approach to AML detection developed specifically for the payment service provider context. The graph analytics module within CyberAid-MONITOR implements transaction network analysis that maps relationships between entities across payment ecosystems enabling the detection of structuring smurfing and layering patterns that would not be visible through transaction-level analysis alone. This capability is complemented by the log analysis and alert correlation functions of CyberAid-MONITOR which integrate authentication events and API access patterns with transaction data to identify coordinated activities across accounts that may indicate organised money laundering operations.

CyberAid-REPORT provides support for the suspicious activity reporting obligations arising under 4AMLD generating investigation narratives and documentation packages from detected suspicious patterns that meet the evidentiary standards required for submission to Financial Intelligence Units. The DIOSCURI digital twin environment further enables investigators to simulate and trace transaction flows in a safe environment supporting hypothesis testing and evidence preservation without disrupting production systems.

These capabilities are validated through Pilot 2 of the CyberAid project which is conducted with OTE/COSMOTE Group Payments and targets a measurable reduction in false positive AML alerts alongside improved detection accuracy for simulated money laundering scenarios. The pilot deployment provides a realistic operational context in which to validate CyberAid's 4AMLD compliance capabilities against the specific transaction volumes and payment patterns of a licensed electronic money institution.

4.1.6 EU AI Act

The EU AI Act established under Regulation (EU) 2024/1689 [29] introduces a comprehensive risk-based governance framework for artificial intelligence systems deployed within the European Union. For CyberAid whose core innovation is precisely the development and deployment of an LLM-based agentic infrastructure for financial cybersecurity the AI Act represents a regulatory instrument whose requirements must be integrated into the platform's design from the outset rather than addressed retrospectively.

The AI Act's risk classification framework categorises AI systems according to the nature of their application and the potential harm their failure could cause. Systems used in the management of critical infrastructure including financial market infrastructure are classified as high-risk under Annex III of the regulation where their failure or malfunction could cause significant disruption to economic activity or public safety. CyberAid's AI components to the extent that they inform or support security decisions affecting financial critical infrastructure are designed to meet the requirements applicable to high-risk AI systems providing a conservative compliance baseline that can be adjusted as supervisory guidance on classification develops.

Among the requirements applicable to high-risk AI systems the most operationally significant for CyberAid's design is the human oversight obligation under Article 14. The regulation requires that high-risk AI systems be designed in such a way that natural persons can effectively oversee their functioning intervene when necessary and override automated decisions. CyberAid addresses this through the human-in-the-loop architecture of the LLM orchestration layer in which all security recommendations of material significance are presented to security analysts through a natural language interface before automated actions are executed. This design ensures that the AI system augments rather than supplants expert judgment in high-stakes security decisions preserving the contextual understanding and professional accountability that automated systems cannot provide.

The AI Act's transparency and explainability requirements under Article 13 require that users of high-risk AI systems receive sufficient information to interpret outputs and understand their limitations. CyberAid's LLM agents are designed to accompany their recommendations with reasoning explanations and confidence indicators drawing on explainable AI methods developed in prior EU-funded research. The technical documentation obligations under Articles 11 and 12 require comprehensive logging of AI system

behaviour which is addressed through CyberAid's observability stack providing continuous monitoring of LLM performance and decision audit trails.

The regulation's data governance requirements under Article 10 impose quality standards on the data used to train validate and test AI systems including requirements to examine data for potential biases that could affect performance. CyberAid's use of open-source European AI models and its reliance on curated financial security knowledge bases rather than general-purpose commercial AI services reflects a deliberate design choice to maintain control over data governance and to reduce dependency on third-country technology providers consistent with the EU's strategic objectives in this area.

CyberAid will implement a conformity assessment pathway for its AI components based on the self-assessment procedures available under the AI Act for high-risk systems not subject to third-party assessment requirements. This process will be supported by the AI Act compliance framework developed in Deliverable D1.5 and will be validated through the pilot deployments in WP6 which will generate operational evidence of the platform's compliance posture under realistic conditions. Table 5 summarises how the main EU regulatory instruments analysed in this section relate to CyberAid capabilities and expected project responses.

Regulation	Primary Obligations	CyberAid Response
DORA (EU) 2022/2554 [8]	ICT risk management, incident reporting, TLPT, third-party risk	CyberAid-MONITOR and CyberAid-REPORT address incident classification and reporting timelines; CyberAid-PROACTIVE supports TLPT and the annual resilience testing program
NIS2 (EU) 2022/2555 [23]	Security measures, supply chain security, management accountability, incident notification	CyberAid-BLUEPRINTS covers NIS2 control domains; CyberAid-MONITOR addresses supply chain behavioral monitoring; dual-trigger notification supported in CyberAid-REPORT
GDPR (EU) 2016/679 [24]	Data minimization, privacy by design, DPIA, breach notification	PySyft federated analytics ensures data never leaves originating institution; configurable retention policies; DPIA documentation produced per pilot deployment
PSD2 (EU) 2015/2366 [26][27]	Operational risk reporting, SCA, open banking API security	CyberAid-PROACTIVE monitors API traffic and OAuth flows; CyberAid-MONITOR supports fraud rate monitoring for SCA exemption eligibility
4AMLD (EU) 2015/849 [28]	Transaction monitoring, suspicious activity reporting, customer due diligence	CyberAid-MONITOR graph analytics detects structuring and layering; CyberAid-REPORT generates SAR documentation; validated through Pilot 2 (OTE/COSMOTE)
EU AI Act (EU) 2024/1689 [29]	High-risk AI obligations, human oversight, transparency, conformity assessment	HITL architecture across CyberAid-LLM; explainability via reasoning outputs and confidence scores; self-assessment conformity pathway supported by D1.5

Table 5: Mapping of EU regulatory requirements to CyberAid capabilities

4.2 Relevant Standards and Frameworks

Alongside the mandatory regulatory framework CyberAid's design is informed by a set of internationally recognized technical standards and operational frameworks that provide implementation guidance benchmarking criteria and interoperability specifications. These standards occupy a different legal status from the regulations described in the preceding section: compliance with them is not directly mandated by law but they are extensively referenced in regulatory instruments supervisory guidance and contractual requirements and they provide the practical implementation layer through which regulatory obligations are given technical expression.

4.2.1 ISO/IEC 27001 and 27035

ISO/IEC 27001:2022 [30] defines the requirements for establishing implementing maintaining and continually improving an Information Security Management System. It provides the structural framework within which CyberAid-BLUEPRINTS reference architectures are designed ensuring that the security controls specified across its annex address the full range of security domains relevant to financial institutions. ISO 27001 is extensively referenced in DORA's implementing technical standards and in supervisory guidance from the European Banking Authority and certification against the standard provides financial institutions with a documented and externally verified basis for demonstrating compliance with key ICT risk management obligations.

ISO/IEC 27035 [31] provides guidance on information security incident management establishing a structured approach to the detection reporting assessment and response to security incidents. The standard's process model covers preparation detection and reporting assessment and decision response and lessons learnt and this structure maps directly onto the workflow architecture of CyberAid-REPORT. The standard's guidance on evidence preservation and forensic investigation is particularly relevant to the digital forensics capabilities developed in WP5 which must produce outputs meeting the evidentiary requirements of both regulatory reporting and where applicable legal proceedings.

4.2.2 NIST Cybersecurity Framework (CSF)

The NIST Cybersecurity Framework [32] updated to version 2.0 in February 2024 organises cybersecurity activities across six core functions: Govern Identify Protect Detect Respond and Recover. Although originating from the United States standards environment NIST CSF has achieved broad international adoption and is explicitly referenced in DORA's implementing technical standards and in the ECB's Cyber Resilience Oversight Expectations framework. It provides a common language for describing cybersecurity posture that facilitates communication between technical teams management bodies and supervisory authorities.

CyberAid maps to the NIST CSF 2.0 functions across its component architecture. The Govern function introduced as a new top-level function in CSF 2.0 is addressed through the policy frameworks and governance templates within CyberAid-BLUEPRINTS. The Identify function is served by CyberAid-PROACTIVE's asset dependency modelling and vulnerability identification capabilities. The Protect function is addressed through the combination of automated penetration testing quantum-resistant cryptographic methods and secure coding capabilities. Detection is provided through CyberAid-MONITOR while Response and Recovery map to CyberAid-REPORT and the incident management workflows of WP5. This comprehensive coverage of all six CSF 2.0 functions provides the validation baseline referenced in KPI-1.1 which targets 100 percent lifecycle coverage verified against NIST CSF.

4.2.3 TIBER-EU Framework

The TIBER-EU framework [33] developed by the European Central Bank and adopted by national competent authorities across the EU provides a standardised methodology for conducting intelligence-led red team testing of financial entities' live production systems. It is the primary reference framework for the threat-led penetration testing obligations introduced by Article 26 of DORA and national implementations of TIBER have been adopted by financial supervisors in most EU Member States.

TIBER-EU divides the testing process into three phases: preparation testing and closure. The preparation phase involves the selection of a threat intelligence provider and red team the definition of the testing scope covering critical functions and supporting systems and the production of a targeted threat intelligence report that informs the attack scenarios to be executed. The testing phase involves the red team conducting realistic attack simulations against the defined scope with a blue team responding without prior knowledge of the exercise. The closure phase involves a comprehensive review of findings remediation planning and regulatory attestation.

CyberAid's proactive security capabilities are designed to complement rather than replicate the TIBER-EU process. The XENON red-teaming framework and DIOSCURI digital twin environment enable financial institutions to conduct extensive pre-production testing before committing to the resource-intensive and formally supervised TIBER-EU engagement improving the quality of findings and the institution's readiness to respond during the formal exercise. CyberAid's threat intelligence integration drawing on STIX/TAXII feeds and financial sector threat sources also provides the intelligence foundation that underpins TIBER-EU's intelligence-led approach ensuring that attack scenarios reflect current and financially-relevant threat actor behaviour.

4.2.4 OWASP Financial API Security Guidelines

The Open Web Application Security Project [34] produces a widely adopted set of security guidelines for web applications and APIs including the API Security Top 10 and the Application Security Verification Standard. In the financial sector context these guidelines are particularly relevant given the proliferation of open banking APIs introduced by PSD2 and the API-driven architecture of modern financial technology platforms. OWASP guidelines are referenced in PCI-DSS secure development lifecycle requirements and in the EBA's guidelines on ICT and security risk management giving them de facto regulatory relevance beyond their status as voluntary standards [35][22].

CyberAid incorporates OWASP guidelines across several components. The DevSecOps code analysis agent developed in WP3 uses OWASP ASVS as its primary reference framework for identifying vulnerabilities in financial application code with detection rules calibrated to the patterns most prevalent in payment and banking software. CyberAid-PROACTIVE's attack generation capabilities include test cases targeting the most commonly exploited API vulnerabilities in open banking environments including broken object level authorisation and implementation weaknesses in OAuth 2.0 flows. The CyberAid Learning Centre incorporates OWASP-based secure coding training specifically designed for developers working on financial applications supporting the capacity building objectives of the project alongside its technical deployment goals.

4.2.5 MITRE ATT&CK and STIX/TAXII

Two further standards underpin CyberAid's threat detection and intelligence-sharing capabilities. The MITRE ATT&CK framework [19], and in particular its coverage of adversary techniques relevant to the financial sector, provides the common taxonomy used to classify and correlate observed activity within CyberAid-MONITOR and CyberAid-REPORT, informing both the Wazuh detection rule mappings and the normalisation of cyber threat intelligence feeds. Complementing this, the STIX and TAXII 2.1 specifications [17][18] provide the structured data format and exchange protocol through which threat intelligence is shared between CyberAid and external sources, supporting the information-sharing arrangements encouraged under Article 45 of DORA. Table 6 summarises how the main standards and operational frameworks considered in this section support the implementation, validation and regulatory alignment of CyberAid components.

Standard / Framework	Applicable CyberAid Components	Regulatory Relevance
ISO/IEC 27001:2022 [30]	CyberAid-BLUEPRINTS, WP3, WP4, WP5	Referenced in DORA ITS and EBA ICT Risk Guidelines; supports ISO 27001 certification pathways for adopting institutions
ISO/IEC 27035 [31]	CyberAid-REPORT, WP5 forensics	Incident management process model aligned with DORA incident reporting obligations and digital forensics evidentiary standards
NIST CSF 2.0 [32]	All CyberAid components	Referenced in DORA ITS and ECB CROE; provides the KPI-1.1 validation baseline covering all six core cybersecurity functions

TIBER-EU [33]	CyberAId-PROACTIVE, Pilot deployments	Primary reference for DORA Article 26 TLPT obligations; pre-TIBER testing supported by XENON and DIOSCURI
OWASP ASVS / API Top 10 [34][35]	CyberAId-PROACTIVE, DevSecOps agent	Referenced in PCI-DSS SDLC requirements and EBA ICT guidelines; governs API security testing for open banking deployments
MITRE ATT&CK (Financial)[19]	CyberAId-MONITOR, CyberAId-REPORT	Threat classification and correlation standard used in Wazuh rule mapping and CTI feed normalisation
STIX / TAXII 2.1[17][18]	CyberAId-MONITOR, CyberAId-REPORT	Threat intelligence exchange standard supporting DORA Article 45 information sharing arrangements

Table 6: Mapping of relevant standards and frameworks to CyberAId components

5. CYBERAID REQUIREMENTS CATALOGUE

This section presents the initial CyberAid requirements catalogue developed in the context of D2.1. The catalogue consolidates requirements collected from project partners, technical work packages, pilot needs, relevant standards and regulations, financial-sector cybersecurity challenges, and the CyberAid Grant Agreement. The purpose of the catalogue is to provide a common baseline for technical design, architecture definition, cybersecurity blueprints, component development, integration activities, and subsequent validation work.

The requirements included in this section are considered as the initial baseline of the project. They reflect the current understanding of the CyberAid scope, the financial-sector context, the expected project capabilities and the needs expressed by partners during the requirements elicitation process. As the project progresses, some requirements may be refined, merged, reprioritised or updated based on the evolution of the architecture, pilot scenarios, technical implementation and validation activities.

Due to the size and variety of requirements, the requirements are not presented as a single large table. Instead, the catalogue is split into different matrices, each one focusing on a specific view of the requirements. This approach improves readability and makes it easier to analyse the requirements from different perspectives, while maintaining traceability to the consolidated requirements repository.

The matrices presented in the following subsections include the functional requirements, non-functional requirements, regulatory and compliance-related requirements, pilot and stakeholder requirements, prioritisation information and traceability mappings. Each matrix is linked to the same underlying requirements baseline, using the requirement identifier as the common reference point.

The requirements are organised according to the main CyberAid capability areas: CyberAid-BLUEPRINTS, CyberAid-PROACTIVE, CyberAid-MONITOR and CyberAid-REPORT. This classification follows the project structure and helps connect the requirements with the components and technical work that will use them. Within each area, requirements are further separated into functional and non-functional requirements. Table 7 provides a compact overview of the initial CyberAid requirements catalogue. The table lists the requirement identifiers, titles and MoSCoW priority values, providing a high-level view before the detailed functional and non-functional matrices are presented in the following subsections.

5.1 Overview of CyberAid Requirements

Req ID	Requirement Title	P	Req ID	Requirement Title	P
R-BLP-01	Finance Security Knowledge Model (SKM) core entities & relations	MH	R-MON-30	Threat Intelligence Agent	MH
R-BLP-02	Standardized ontologies & taxonomies	MH	R-MON-31	Automated Failover	MH
R-BLP-03	CTI Ingestion, Normalization and External Threat Intelligence Integration	MH	R-MON-32	Multi-Agent Orchestration Pipeline	MH
R-BLP-04	Semantic indexing for LLM reasoning	MH	R-PRO-09	Vulnerability Assessment Agent	MH
R-BLP-05	CTI governance: integrity, access control, and auditability	MH	R-PRO-10	DevSecOps & Code Analysis Agent	SH
R-BLP-06	Privacy-by-design for CTI enrichment & sharing	MH	R-PRO-11	LLM Input Protection	MH
R-BLP-07	Specialized Agent — Dual-Mode Interface Contract	MH	R-PRO-12	API Authentication & Encryption	MH

R-BLP-08	Partitioned Vector Store per Agent	MH	R-PRO-13	Privacy-Preserving Data Handling	MH
R-BLP-09	Concurrent Agent Operation	MH	R-PRO-14	Q2FA Protocol, API, Simulation and Pilot Integration	SH
R-BLP-10	Governed Tool Adapter Layer for All External Services	MH	R-PRO-15	Quantum-Resistant Cryptography Assessment and PoC	SH
R-BLP-11	Configuration-Driven Threat Ontology	MH	R-PRO-16	Privacy-Preserving Federated Analytics for Sensitive Data	MH
R-BLP-12	Hybrid Retrieval (Semantic + Keyword)	MH	R-PRO-17	Digital Twin Environment Provisioning	MH
R-BLP-13	External Knowledge Source Integration	MH	R-PRO-18	Service Replication Fidelity	MH
R-BLP-14	On-Premises Deployment	MH	R-PRO-19	Network Traffic Simulation in Twin	SH
R-BLP-15	Historical Data Ingestion	CH	R-PRO-20	Test Lifecycle Orchestration	MH
R-BLP-16	Data Leakage Prevention at Agent Output	MH	R-PRO-21	Digital Twin Isolation & Security	MH
R-MON-01	Continuous Network Monitoring for Financial Infrastructure	MH	R-PRO-22	Common data and model structure	MH
R-MON-02	Network Monitoring Performance Overhead	SH	R-PRO-23	Automated Attack Surface Discovery	MH
R-MON-03	Financial Application-Specific Monitoring Profiles	MH	R-PRO-24	AI-Driven Exploit Chain Generation	MH
R-MON-04	Real-Time Connection Tracing and Security Event Detection	MH	R-PRO-25	Continuous & Scheduled Pentest Execution	SH
R-MON-05	Network-Application Behaviour Correlation	MH	R-PRO-26	Pentest Finding Normalisation & Risk Integration	MH
R-MON-06	Behavioural Analysis Agent — UEBA	MH	R-PRO-27	Regulatory Compliance Mapping of Pentest Findings	SH
R-MON-07	Response Latency — Critical Events	MH	R-REP-01	Forensic Evidence Collection with Chain of Custody	MH
R-MON-08	Multi-Vector Threat Coverage	MH	R-REP-02	Financial-Specific Forensic Workflows and Timeline Reconstruction	MH
R-MON-09	Real-Time Detection	MH	R-REP-03	Threat Intelligence Integration for Forensic Contextualization	MH
R-MON-10	Sender / Entity Anomaly Correlation	CH	R-REP-04	Financial Specific Root Cause Analysis	MH
R-MON-11	API Bridge Security Monitoring	MH	R-REP-05	Privacy-Preserving Forensic Investigation	MH

R-MON-12	Market Data Latency Monitoring	MH	R-REP-06	Result Aggregation, Corroboration and Conflict Resolution	MH
R-MON-13	Execution Latency Monitoring	MH	R-REP-07	Compliance Verification Agent	MH
R-MON-14	Market Data Feed Integrity	MH	R-REP-08	Incident Response Agent	MH
R-MON-15	Multi-Modal Impersonation Detection	CH	R-REP-09	Forensic Analysis Agent	MH
R-MON-16	Secure Timely Notification	CH	R-REP-10	Hallucination Mitigation	MH
R-MON-17	Geographic Locality Analysis	CH	R-REP-11	Explainable AI Output for Security and Identity Decisions	MH
R-MON-18	Deployment-Time Context Configuration	SH	R-REP-12	Actionable Response Recommendations	MH
R-MON-19	Algorithm Repository Monitoring	MH	R-REP-13	Full Audit Trail and Incident Logging	MH
R-MON-20	Performance Impact Control	MH	R-REP-14	REST API and Analyst Dashboard	MH
R-MON-21	AI Risk Modelling	SH	R-REP-15	Regulatory Reporting	SH
R-MON-22	Common monitoring APIs and logs submission	SH	R-REP-16	Executive, KPI and Risk Dashboard	SH
R-MON-23	Graph-based modelling of financial entities & transactions	MH	R-REP-17	Risk-Weighted Confidence Scoring	MH
R-MON-24	Neo4j + Cypher pattern matching capability	MH	R-REP-18	Holistic Multi-Source Risk Aggregation	MH
R-MON-25	Specialized algorithms (NetworkX/Python) for fraud detection	MH	R-REP-19	DORA-Aligned ICT Risk Metrics & Reporting	MH
R-MON-26	Analyst-facing graph visualization for investigations	SH	R-REP-20	Remediation Workflow & Treatment Tracking	MH
R-MON-27	Integration of graph intelligence with risk/LLM layers	MH	R-REP-21	Risk Trend Analysis & Predictive Risk Indicators	CH
R-MON-28	Real-time readiness: scalability & latency constraints	MH	R-REP-22	CRA — Event Triage & Classification	MH
R-MON-29	CRA — Scheduled & Periodic Analysis	SH	R-REP-23	CRA — Multi-Agent Parallel Dispatch	MH

Table 7: Requirement Catalogue

5.2 Functional Requirements

This subsection presents the functional requirements of the CyberAid framework. The requirements are grouped according to the main CyberAid capability areas and describe the capabilities expected from each area. Additional metadata, including source, priority, validation means and WP/task mapping, is maintained in the consolidated requirements repository and reflected in the annex where applicable.

5.2.1 Blueprint Requirements

Table 8 presents the functional requirements related to CyberAid-BLUEPRINTS. These requirements focus on the security knowledge model, CTI integration, semantic indexing, agent deployment patterns, external knowledge sources and blueprint-related architectural capabilities.

Req ID	Source Type	Requirement Title	Requirement Description
R-BLP-01	Other	Finance Security Knowledge Model (SKM) core entities & relations	CyberAid shall implement a finance-focused SKM capturing relationships among threats, vulnerabilities, assets, controls, and regulatory requirements, enabling consistent interpretation and analysis across the platform.
R-BLP-02	Other	Standardized ontologies & taxonomies	CyberAid shall define and publish ontologies/taxonomies that standardize security concepts across CyberAid components and support interoperability with industry schemas.
R-BLP-03	Other	CTI Ingestion, Normalization and External Threat Intelligence Integration	CyberAid shall specify and implement CTI integration mechanisms to ingest, normalize, correlate and preserve provenance/confidence for intelligence from financial ISAC feeds, STIX/TAXII sources, open-source repositories and other external threat intelligence providers.
R-BLP-04	Other	Semantic indexing for LLM reasoning	CyberAid shall support semantic indexing grounded to SKM entities to enable LLM-based reasoning over both structured CTI and unstructured sources (bulletins, reports), including extraction patterns for relevant threat knowledge.
R-BLP-07	Architecture	Specialized Agent — Dual-Mode Interface Contract	Every specialized domain agent shall expose a standardised dual-mode interface: (a) orchestrated mode — invokable by the Central Router Agent (CRA), with Security Context enforcement and append-only audit on every invocation; and (b) standalone mode — invokable directly via CLI or REST API. MCP Server wrapping is supported as an optional external-facing adapter to expose agents to third-party consumers or standalone microservice deployment contexts where the platform runtime is not present.
R-BLP-08	Architecture	Partitioned Vector Store per Agent	The knowledge infrastructure shall maintain dedicated vector store partitions per agent domain (threat intel, vulnerability, compliance, behavioural, etc.) plus a shared/global partition, with all retrieval scoped by Security Context.
R-BLP-10	Architecture	Governed Tool Adapter Layer for All External Services	Domain agents shall access all Layer 2 security services (Wazuh, OLISTIC, XÉNON, DIOSCURI, CTI feeds) through a governed Tool Adapter Layer that enforces Security Context checks, JSON schema argument validation, and append-only audit logging on every external call, regardless of the

			underlying transport. Adapters are implemented as typed interfaces; the platform ships native typed adapters as the default. MCP Tool Server wrapping is supported as an optional integration pattern for services that expose an MCP-compatible interface or for external consumption of the adapter surface.
R-BLP-11	Architecture	Configuration-Driven Threat Ontology	Attack types, agent-to-attack mappings, priority weights, and thresholds shall be defined in configuration, extensible without code changes to the core pipeline.
R-BLP-12	Architecture	Hybrid Retrieval (Semantic + Keyword)	The RAG pipeline shall support hybrid retrieval combining semantic vector similarity search with keyword-based filtering, plus re-ranking and deduplication of retrieved chunks.
R-BLP-13	Standard	External Knowledge Source Integration	The knowledge infrastructure shall ingest and index content from CVE/NVD, MITRE ATT&CK, OWASP, STIX/TAXII CTI feeds, and relevant financial regulatory frameworks (DORA, GDPR, PCI DSS).
R-BLP-15	Pilot	Historical Data Ingestion	The system should support the ingestion of historical "suspicious activity" logs from KM3 to train the local AI model.

Table 8: Blueprint Functional Requirements

5.2.2 Monitor Requirements

Table 9 presents the functional requirements related to CyberAid-MONITOR. These requirements cover continuous monitoring, anomaly detection, behavioural analysis, API and market data monitoring, graph-based fraud detection, threat intelligence analysis and multi-agent monitoring orchestration.

Req ID	Source Type	Requirement Title	Requirement Description
R-MON-01	Other	Continuous Network Monitoring for Financial Infrastructure	The platform shall implement continuous network monitoring to capture network traffic and application behaviour across financial infrastructure without requiring modifications to the monitored systems.
R-MON-03	Other	Financial Application-Specific Monitoring Profiles	The platform shall provide specialized monitoring profiles for common financial applications and services, including payment processors, banking platforms, trading systems, and authentication services, ensuring visibility across the financial technology stack.
R-MON-04	Other	Real-Time Connection Tracing and Security Event Detection	The platform shall provide real-time visibility into all active network connections and their associated processes, enabling immediate detection of unauthorized communications and data exfiltration attempts in financial environments. The solution shall include financial-specific detection rules to monitor and detect security-relevant events, including unusual data access patterns and privilege escalation attempts.

R-MON-05	Other	Network-Application Behaviour Correlation	The platform shall correlate network-level activities with application behaviours, creating a holistic view of system operations that enhances attack detection capabilities.
R-MON-06	Architecture	Behavioural Analysis Agent — UEBA	The Behavioural Analysis Agent shall perform user and entity behaviour analytics (UEBA), detecting insider threats and anomalous behaviour patterns using deviation from established baselines.
R-MON-08	Pilot	Multi-Vector Threat Coverage	The system shall support detection across multiple input vectors relevant to financial services, including email-based attacks, log-based anomalies, and behavioural deviations. A domain-specific threat ontology shall define attack types, detection agents, and trigger rules per vector.
R-MON-10	Pilot	Sender / Entity Anomaly Correlation	The system shall correlate contextual metadata (e.g., geographic origin, timing, communication patterns) with known entity profiles to detect behavioural deviations.
R-MON-11	Pilot	API Bridge Security Monitoring	Build and monitor API traffic between TradeStation and Interactive Brokers for suspicious patterns or cyber threats.
R-MON-12	Pilot	Market Data Latency Monitoring	Continuously monitor latency of market data feeds from IB and Tradestation.
R-MON-13	Pilot	Execution Latency Monitoring	Monitor latency between order submission and execution confirmation from Interactive Brokers.
R-MON-14	Pilot	Market Data Feed Integrity	Detect anomalies or irregular patterns in TradeStation/IB data feeds.
R-MON-15	Pilot	Multi-Modal Impersonation Detection	The system must detect suspicious behaviour in client emails/voice transcripts/ and other ways
R-MON-16	Pilot	Secure Timely Notification	Incident alerts shall be delivered to the appropriate compliance or operational stakeholder through approved secure channels within a configurable SLA after detection.
R-MON-17	Pilot	Geographic Locality Analysis	The system shall correlate the geographic origin of incoming communications, where lawful and technically available through metadata, with expected client context to flag unusual locality-based deviations for analyst review.
R-MON-18	Architecture	Deployment-Time Context Configuration	The system shall accept deployment-specific context (e.g., employee directories, known domains, trusted accounts, organisational hierarchy) to improve detection accuracy. Agents shall degrade gracefully when context is absent.
R-MON-19	Pilot	Algorithm Repository Monitoring	Monitor EasyLanguage algorithm repositories for unauthorized access or modification using eBPF and AI anomaly detection.

R-MON-21	Pilot	AI Risk Modelling	Provide real-time AI-powered threat scoring for anomalies in algorithms, market data feeds, or API traffic.
R-MON-22	Technology	Common monitoring APIs and logs submission	Given that CyberAid may need logging across diverse systems and distributed sites, a common logging mechanism should be in place, to hold such information. A unified interface should exist in order to aid all the individual systems to report and store logs from any part of the pipeline.
R-MON-23	Other	Graph-based modelling of financial entities & transactions	CyberAid shall implement graph-based modelling of financial relationships, transactions, and entity behaviours to identify suspicious transaction networks and patterns not captured by rule-based systems.
R-MON-24	Other	Neo4j + Cypher pattern matching capability	The graph analytics engine shall use Neo4j as a core graph store and provide Cypher-based pattern matching for fraud typologies (e.g., rings, rapid fund movement, layering/structuring motifs) over temporal transaction graphs.
R-MON-25	Other	Specialized algorithms (NetworkX/Python) for fraud detection	CyberAid shall implement specialized graph analytics algorithms (using NetworkX/Python graph libraries) for mapping transaction flows, account relationships, and temporal patterns to detect coordinated fraudulent behaviours and anomalies.
R-MON-26	Other	Analyst-facing graph visualization for investigations	CyberAid shall provide financial-specific graph visualizations enabling analysts to explore suspicious networks (filtering by time windows, entity types, risk scores, and path tracing) to support investigations.
R-MON-27	Other	Integration of graph intelligence with risk/LLM layers	Graph-derived intelligence shall be exportable via a defined interface to (i) the risk assessment engine (OLISTIC) and (ii) the LLM orchestration layer for contextual enrichment and explanation.
R-MON-29	Architecture	CRA — Scheduled & Periodic Analysis	The CRA shall support time-triggered orchestration patterns (e.g., nightly vulnerability sweeps, weekly compliance posture summaries) in addition to event-driven workflows.
R-MON-30	Architecture	Threat Intelligence Agent	The Threat Intelligence Agent shall perform IoC analysis, TTP mapping to MITRE ATT&CK, threat actor attribution, and campaign correlation by querying CTI feeds and the dedicated knowledge partition.
R-MON-31	Pilot	Automated Failover	Automatically switch to backup market data feeds when anomalies or latency thresholds are exceeded.
R-MON-32	Architecture	Multi-Agent Orchestration Pipeline	The framework shall implement a multi-agent detection pipeline using a graph-based orchestration framework (LangGraph). Agents shall execute in parallel where possible, with results fused by an orchestrator using weighted confidence scoring.

Table 9: Monitor Functional Requirements

5.2.3 Proactive Requirements

Table 10 presents the functional requirements related to CyberAId-PROACTIVE. These requirements address vulnerability assessment, DevSecOps analysis, LLM input protection, privacy-preserving analytics, digital twin provisioning, attack surface discovery and AI-driven penetration testing.

Req ID	Source Type	Requirement Title	Requirement Description
R-PRO-09	Architecture	Vulnerability Assessment Agent	The Vulnerability Assessment Agent shall perform risk scoring against CVE/NVD and MITRE databases, assess exploitability in the financial context, and generate prioritized remediation plans with rationale.
R-PRO-10	Architecture	DevSecOps & Code Analysis Agent	The DevSecOps Agent shall perform static and dynamic code analysis, identify security anti-patterns, and verify SDLC compliance for financial software artifacts.
R-PRO-11	Architecture	LLM Input Protection	CyberAId shall protect LLM-based components against malicious or unsafe inputs by combining prompt injection detection, input validation, sanitisation, and policy-based blocking before requests are processed by domain agents.
R-PRO-12	Pilot	API Authentication & Encryption	External API integrations used by CyberAId shall enforce authenticated and encrypted communication, using approved mechanisms such as TLS and token- or certificate-based authentication.
R-PRO-14	Pilot	Q2FA Protocol, API, Simulation and Pilot Integration	CyberAId shall provide a consolidated Quantum Two-Factor Authentication (Q2FA) capability covering the protocol definition, API data model, API adaptation, simulation environment, simulation tests, and pilot application integration for the selected Q2FA use case.
R-PRO-15	Architecture	Quantum-Resistant Cryptography Assessment and PoC	CyberAId shall assess quantum-resistant cryptography options for prioritised secure connections or transactions and provide a proof-of-concept evaluation showing feasibility, integration constraints and recommended deployment approach.
R-PRO-16	Regulation	Privacy-Preserving Federated Analytics for Sensitive Data	CyberAId shall support privacy-preserving analytics and federated training patterns for sensitive financial data, enabling model improvement or collaborative analytics without centralising raw personal or confidential data.
R-PRO-17	Other	Digital Twin Environment Provisioning	CyberAId shall provide an automated mechanism to instantiate replicated digital twin environments that mirror target production service topologies, configurations, and data flows for safe, isolated testing without impacting live systems.
R-PRO-18	Other	Service Replication Fidelity	Digital twin environments shall faithfully reproduce the behaviour of deployed services, including API interfaces, authentication flows, data schemas, and inter-service dependencies, to within a defined fidelity tolerance, enabling representative security and functional testing.

R-PRO-19	Other	Network Traffic Simulation in Twin	The digital twin platform shall support configurable network traffic simulation, including synthetic load generation, latency injection, packet loss, and replay of captured production traffic traces, to enable realistic security and performance testing scenarios.
R-PRO-20	Standard	Test Lifecycle Orchestration	CyberAId shall provide an orchestration layer for the full test lifecycle within digital twin environments, covering test plan execution, result capture, environment teardown, and auditability of test runs, with traceability from test case to requirement.
R-PRO-22	Technology	Common data and model structure	In the case of federated learning across different entities, in order to jointly train AI models on a common task, the data formats and selected inputs for the given model need to agree
R-PRO-23	Regulation	Automated Attack Surface Discovery	The pentesting module shall automatically enumerate the attack surface of CyberAId-monitored financial services, including exposed APIs, network services, authentication endpoints, and third-party integrations, producing a continuously updated attack surface inventory for each pentest cycle.
R-PRO-24	Other	AI-Driven Exploit Chain Generation	The pentesting module shall leverage AI reasoning (hex-clad-ai) to generate multi-step exploit chains by chaining discovered vulnerabilities, misconfigurations, and credential weaknesses into realistic attack scenarios representative of advanced persistent threats targeting financial infrastructure.
R-PRO-25	Other	Continuous & Scheduled Pentest Execution	The pentesting module shall support both on-demand and scheduled pentest execution against digital twin environments, with configurable scope, intensity, and technique profiles, and shall integrate with the CyberAId CI/CD pipeline to trigger security regression tests on each deployment.
R-PRO-26	Standard	Pentest Finding Normalisation & Risk Integration	Pentest findings shall be normalised to a standard vulnerability taxonomy (CVE/CWE/CVSS) and automatically ingested into the CyberAId risk management engine (OLISTIC) to enrich the holistic risk score and trigger remediation workflows, with bidirectional traceability between findings and risk records.
R-PRO-27	Regulation	Regulatory Compliance Mapping of Pentest Findings	The pentesting module shall automatically map discovered vulnerabilities and findings to applicable regulatory controls (DORA, NIS2, GDPR, PSD2, AI Act) and generate compliance gap reports that identify which regulatory obligations are at risk, supporting evidence-based remediation prioritisation.

Table 10: Proactive Functional Requirements

5.2.4 Report Requirements

Table 11 presents the functional requirements related to CyberAid-REPORT. These requirements focus on forensic evidence collection, incident response, compliance verification, explainable AI outputs, audit trails, dashboards, risk aggregation and regulatory reporting.

Req ID	Source Type	Requirement Title	Requirement Description
R-REP-01	Other	Forensic Evidence Collection with Chain of Custody	The system will implement evidence collection mechanisms designed to preserve volatile data from financial systems while maintaining chain of custody and forensic soundness.
R-REP-02	Other	Financial-Specific Forensic Workflows and Timeline Reconstruction	The platform shall provide forensic analysis workflows for reconstructing security incidents in financial environments, combining multiple data sources into coherent incident narratives.
R-REP-03	Other	Threat Intelligence Integration for Forensic Contextualization	The system shall integrate with threat intelligence repositories in order to contextualize forensic findings within the broader financial threat landscape, identifying potential threat actors and their tactics, techniques, and procedures.
R-REP-04	Other	Financial Specific Root Cause Analysis	The platform shall support root cause analysis specifically adapted for financial environments, identifying underlying vulnerabilities and recommending remediation actions.
R-REP-06	Architecture	Result Aggregation, Corroboration and Conflict Resolution	The CyberAid orchestration layer shall aggregate multi-agent outputs, corroborate findings across sources, resolve conflicts, and produce a unified explanation and prioritised result for downstream reporting and response.
R-REP-07	Legislation	Compliance Verification Agent	The Compliance Verification Agent shall map identified security events and posture assessments against DORA, GDPR, and PCI DSS regulatory frameworks, flagging gaps and generating compliance evidence artifacts.
R-REP-08	Architecture	Incident Response Agent	The Incident Response Agent shall generate containment strategies, recommend playbook activation, and support recovery planning for confirmed incidents.
R-REP-09	Architecture	Forensic Analysis Agent	The Forensic Analysis Agent shall perform timeline reconstruction, support evidence preservation procedures, and produce forensic narratives meeting evidentiary standards for financial-sector regulatory obligations.
R-REP-12	Pilot	Actionable Response Recommendations	Detection outputs shall be translated into prioritised, role-appropriate actionable tasks for end users (e.g., block, quarantine, escalate, investigate).
R-REP-14	Architecture	REST API and Analyst Dashboard	The framework shall expose a REST API for data submission and provide a web dashboard for analysts to inspect findings, filter by severity, and review explanations.
R-REP-16	Pilot	Executive, KPI and Risk Dashboard	CyberAid shall provide role-based dashboards and reports summarising security incidents, activity volumes, false-positive rates, algorithm security status, API health, market/execution latency, alerts, risk posture and remediation status for technical and executive users.

R-REP-17	Pilot	Risk-Weighted Confidence Scoring	The system must provide a "Confidence Score" (0-100%) for identity verification and/or other fraud detection likelihood, allowing staff to quickly assess risk levels.
R-REP-18	Other	Holistic Multi-Source Risk Aggregation	The OLISTIC risk management engine shall aggregate risk signals from all CyberAid modules — CTI feeds, graph analytics, LLM detections, SIEM alerts, and pentest findings — into a unified, continuously updated risk posture model, applying a configurable weighting and aggregation methodology documented in the platform design.
R-REP-20	Standard	Remediation Workflow & Treatment Tracking	OLISTIC shall support a structured risk treatment workflow enabling operators to assign remediation tasks to responsible owners, set deadlines, track progress, record treatment decisions (accept/mitigate/transfer/avoid), and automatically re-evaluate residual risk upon completion, with a full audit trail of all treatment decisions.
R-REP-21	Regulation	Risk Trend Analysis & Predictive Risk Indicators	OLISTIC shall analyse historical risk data to compute risk trend indicators and, where feasible, provide predictive risk signals using ML-based forecasting, alerting operators to deteriorating risk postures before threshold breaches occur, with explainable outputs traceable to contributing risk factors.
R-REP-22	Architecture	CRA — Event Triage & Classification	The CRA shall receive security events from the message queue, perform initial severity assessment and incident categorization, and estimate potential impact before dispatching to specialized agents.
R-REP-23	Architecture	CRA — Multi-Agent Parallel Dispatch	The CRA shall be capable of simultaneously activating multiple specialized agents when an incident requires multi-domain analysis, coordinating parallel execution and collecting partial results as they arrive.

Table 11: Report Functional Requirements

5.3 Non-Functional Requirements

5.3.1 Blueprint Requirements

Table 12 presents the non-functional requirements related to CyberAid-BLUEPRINTS. These requirements address CTI governance, privacy-by-design, concurrent agent operation, on-premises deployment and data leakage prevention.

Req ID	Source Type	Requirement Title	Requirement Description
R-BLP-05	Standard	CTI governance: integrity, access control, and auditability	CTI and SKM datasets shall be protected with role-based access control, integrity protection, and auditable change history (who changed what/when/why) for all model and feed configuration updates.
R-BLP-06	Regulation	Privacy-by-design for CTI enrichment & sharing	When enriching or sharing CTI, CyberAid shall apply data minimization and privacy-by-design safeguards (e.g., pseudonymization/redaction of personal data; purpose limitation) and document confidentiality tiers for bidirectional sharing.

R-BLP-09	Architecture	Concurrent Agent Operation	The system shall support concurrent request handling from both CRA and Reporting Agent to all domain agents, with stateless agent design ensuring isolation between concurrent sessions.
R-BLP-14	Pilot	On-Premises Deployment	The detection framework shall be deployable entirely on premises. No data shall leave the institution's network. Local LLM inference shall be supported as an alternative to cloud APIs.
R-BLP-16	Architecture	Data Leakage Prevention at Agent Output	Every agent output pipeline shall include a Data Leakage Prevention check verifying that no data outside the requesting user's permitted Security Context scope appears in the response — mandatory for multi-tenant deployments.

Table 12: Blueprint Non-Functional Requirements

5.3.2 Monitor Requirements

Table 13 presents the non-functional requirements related to CyberAid-MONITOR. These requirements focus on performance overhead, response latency, real-time detection, monitoring performance impact and scalability constraints.

Req ID	Source Type	Requirement Title	Requirement Description
R-MON-02	Other	Network Monitoring Performance Overhead	The monitoring solution shall not impose performance degradation on production financial workloads.
R-MON-07	Architecture	Response Latency — Critical Events	For CRITICAL severity events, the CRA shall deliver an initial triage assessment within 60 seconds of event receipt and a complete multi-agent synthesized finding within 5 minutes.
R-MON-09	Pilot	Real-Time Detection	Threat identification shall occur in near-real-time at the point of ingestion (email receipt, log event arrival). The detection pipeline shall meet defined latency KPIs per severity level.
R-MON-20	Pilot	Performance Impact Control	Monitoring and analysis components for trading operations shall operate within a defined performance budget, with latency and CPU overhead measured under representative workloads.
R-MON-28	Regulation + General	Real-time readiness: scalability & latency constraints	The graph analytics pipeline shall be engineered for scale (partitioning/parallel processing as needed) to maintain near-real-time fraud pattern detection under peak transaction volumes, with defined performance KPIs and monitored SLOs.

Table 13: Monitor Non-Functional Requirements

5.3.3 Proactive Requirements

Table 14 presents the non-functional requirements related to CyberAid-PROACTIVE. These requirements mainly concern privacy-preserving data handling and isolation of digital twin environments from production systems.

Req ID	Source Type	Requirement Title	Requirement Description
--------	-------------	-------------------	-------------------------

R-PRO-13	Legislation	Privacy-Preserving Data Handling	Personal data processed during email/voice/and other analysis must be encrypted and handled according to the principle of "Data Minimization."
R-PRO-21	Regulation	Digital Twin Isolation & Security	Digital twin environments shall be strictly isolated from production systems and networks, with enforced access controls, secrets management for any injected credentials, and automated sanitisation of any production data used to seed test environments, ensuring no risk of data leakage or cross-environment interference.

Table 14: Proactive Non-Functional Requirements

5.3.4 Report Requirements

Table 15 presents the non-functional requirements related to CyberAid-REPORT. These requirements focus on privacy-preserving forensic investigation, hallucination mitigation, explainability, auditability, regulatory reporting and DORA-aligned risk metrics.

Req ID	Source Type	Requirement Title	Requirement Description
R-REP-05	Legislation	Privacy-Preserving Forensic Investigation	The forensic toolkit shall implement privacy-preserving investigation techniques while maintaining GDPR compliance.
R-REP-10	Architecture	Hallucination Mitigation	The output processing pipeline shall include hallucination mitigation mechanisms (e.g., evidence grounding checks, source attribution validation) before delivering findings to orchestrators or users.
R-REP-11	Legislation	Explainable AI Output for Security and Identity Decisions	CyberAid shall provide explainable AI outputs for security, fraud and identity-verification decisions, including the evidence used, confidence score or rationale, source citations where available, and auditability suitable for regulatory and operational review.
R-REP-13	Legislation	Full Audit Trail and Incident Logging	Every CyberAid analysis and relevant security incident shall produce a structured, machine-readable audit trail covering routing decisions, agent findings with evidence, scoring breakdowns, stage timings, final verdicts, API/market/execution events where relevant, and tamper-evident preservation for compliance and forensic use.
R-REP-15	Pilot	Regulatory Reporting	Generate BaFin/Bundesbank-compliant reports for incidents affecting trading operations.
R-REP-19	Regulation	DORA-Aligned ICT Risk Metrics & Reporting	OLISTIC shall implement the ICT risk management metrics and reporting structures required by DORA, including risk appetite thresholds, key risk indicators (KRIs), and automated generation of ICT risk reports suitable for submission to competent authorities and for board-level governance.

Table 15: Report Non-Functional Requirements

5.4 Pilot Stakeholder Requirements

5.4.1 Pilot #1 – Client Impersonation Detection

Pilot #1 focuses on Client Impersonation Detection. The pilot addresses the unique security challenges in the complex ecosystem of private banking and asset management. Smaller asset managers frequently relay client orders to larger custodian institutions. This process creates potential security gaps where attackers can insert fraudulent instructions via compromised email accounts or endpoint devices. The objective is to identify these sophisticated impersonation attempts, which often fall below standard verification thresholds (e.g., €50K) and can result in significant aggregated losses while evading traditional security controls and causing reputation damage.

The technical contribution of Pilot #1 lies in the design of an impersonation detection system that leverages email content analysis, federated learning across institutions, adaptive monitoring, and LLM-based contextual reasoning. The system performs deep analysis of email content, structure, and metadata to identify anomalous patterns. It utilizes CyberAid LLM agents with few-shot learning to detect subtle linguistic deviations, unusual attachment types, and suspicious embedded links. To support this, CyberAid establishes a privacy-preserving federated learning network across financial institutions, which enables the development of shared impersonation detection models while keeping actual client data within its originating institution. Based on patterns identified across the institutional network, the system dynamically adjusts detection rules, correlation parameters, and alert thresholds. As a forward-looking measure, the pilot also incorporates a proof-of-concept simulation for Quantum Token 2-factor Authentication to protect connections and interactions between transaction partners.

The successful implementation of Pilot #1 relies on the continuous collaboration among several consortium partners. KM3 acts as the business owner and leads the pilot activities. KM3 Data & Operations supports this by providing vital data assets, such as historical email communications, and ensures operational deployment within their system. On the technical side, INNOV leads the LLM implementation responsible for advanced contextual reasoning, while HUA develops the underlying privacy-preserving federated learning infrastructure. UBI is tasked with providing the overall cybersecurity platform integration to bring the components together. Finally, QUBO undertakes the implementation of the quantum authentication prototype to effectively simulate and demonstrate the advanced use case.

The validation of Pilot #1 will be measured against specific Key Performance Indicators (KPIs) designed to evaluate both detection accuracy and overall operational impact. The core objective is achieving a significantly higher accuracy in identifying simulated fraudulent instructions compared to legacy baselines, while simultaneously attaining a measurable reduction in false alerts through the application of LLM contextual analysis. In terms of the federated model's performance, the pilot aims to ensure that smaller institutions can achieve detection rates comparable to the capabilities demonstrated by larger institutions. Privacy preservation is a critical metric, requiring strict adherence to data minimization and zero-exposure principles during model training, a standard that will be strictly validated by audit. Ultimately, to protect the business flow, the pilot must maintain a seamless client experience by ensuring that enhanced verification of friction for legitimate transactions is kept to an absolute minimum.

Pilot #1 enables smaller financial institutions to benefit from collective intelligence without compromising client data privacy. By implementing differential privacy mechanisms and secure aggregation protocols, the solution significantly enhances detection capabilities for smaller institutions that typically lack the data volume, security resources, and technical capacity to develop robust models independently. Furthermore, the automated adaptation of security monitoring parameters ensures that these smaller institutions benefit from advanced detection capabilities typically reserved for large financial organizations. The expected outcomes include demonstrated cross-institutional value with security improvements for institutions of all sizes. The continuous refinement of detection rules based on cross-institutional learning creates a security monitoring system that stays ahead of evolving attack methodologies. Ultimately, this

secures the critical financial supply chain against sophisticated impersonation attacks and establishes a new standard for collaborative, privacy-preserving cybersecurity in the asset management sector.

5.4.2 Pilot #2 – Payment Processing / iSOC

Pilot #2 focuses on the deployment and validation of CyberAid's AI-powered cybersecurity technologies in the context of OTE's electronic payment service provider operations, with emphasis on Anti-Money Laundering (AML) detection scenarios relevant to Payzy and comparable digital payment environments. The pilot will address the need to identify sophisticated money laundering patterns across digital payment ecosystems, where high transaction volumes, multiple payment channels and complex relationships between users, accounts and merchants may reduce the effectiveness of traditional rule-based AML monitoring. Depending on data availability, internal approvals and privacy/security constraints, the validation may rely on representative pilot data, anonymised or synthetic datasets, sandboxed configurations and/or simulated payment environments rather than direct use of production infrastructure. The objective is to enhance the detection, investigation and prioritisation of suspicious transaction patterns through transaction network analysis, AML-oriented log correlation and LLM-supported alert investigation, while ensuring that the resulting approach remains explainable, auditable and aligned with the operational and regulatory needs of payment service providers.

Technical Contribution

The technical contribution of Pilot #2 lies in the definition and validation of a multi-layered AML detection approach for payment service provider scenarios. The pilot will examine how CyberAid technologies can support the identification of suspicious transaction behaviours by combining transaction network analysis, security log correlation and AI-supported alert investigation. The validation environment and data approach will be defined according to OTE's internal approvals, privacy requirements and technical feasibility, and may rely on synthetic, anonymised, pseudonymised or simulated datasets representing Payzy-relevant payment scenarios.

At the transaction analysis level, the pilot will explore the use of graph-based analytics to represent relationships between users, accounts, merchants, devices and transactions. This capability is expected to support the identification of suspicious AML patterns such as structuring, smurfing, layering or other transaction behaviours that may not be visible through isolated rule-based checks.

At the monitoring and correlation level, CyberAid-MONITOR will be considered for the ingestion and analysis of AML-relevant operational signals, such as transaction logs, authentication events, API access patterns and other payment-service-related indicators, subject to data availability and access constraints. Custom AML-oriented correlation rules may be defined to connect transaction behaviour with technical

events, enabling the pilot to investigate whether suspicious financial activity can be better prioritised when combined with contextual cybersecurity evidence.

At the investigation-support level, the pilot will assess how CyberAid's LLM-based capabilities can assist analysts in understanding and documenting suspicious transaction patterns. The LLM layer may support alert summarisation, contextual explanation, relationship interpretation and preparation of structured investigation narratives, while keeping final assessment and decision-making under human control. The use of DIÓSCURI digital twin capabilities may also support the safe simulation of suspicious transaction scenarios in a controlled environment, enabling testing and validation without requiring exposure of production systems.

Overall, the technical contribution of the OTE pilot is to translate AML monitoring needs from the payment service provider domain into CyberAid technical requirements for graph analytics, monitoring integration, alert correlation, explainability, data protection and controlled validation. The pilot will therefore provide evidence on how CyberAid can support more effective AML detection and investigation workflows while respecting operational, regulatory and confidentiality constraints.

Stakeholder Contributions

OTE will support the pilot by bringing the perspective of a payment service provider and by helping define realistic AML scenarios, stakeholder needs and validation constraints. As the pilot owner, OTE will ensure that the requirements reflect practical challenges faced in digital payment environments, including operational, compliance and security considerations.

OTE will help identify the AML scenarios that are most relevant for the pilot. These may include patterns such as structuring, smurfing, layering, suspicious account relationships, abnormal transaction velocity or other behaviours that are relevant to payment service provider operations.

OTE will also help define the type of data and events that can be used for validation. Depending on internal approvals and privacy or security constraints, the pilot may use synthetic, anonymised, pseudonymised or simulated transaction data, together with representative events such as authentication, API access, device/session or payment-related logs. The intention is to keep the pilot realistic without exposing production systems or sensitive information.

OTE will provide feedback on whether the CyberAid outputs are useful, understandable and auditable from an operational point of view. This includes the outputs of graph-based analysis, alert correlation and LLM-supported investigation. In all cases, final interpretation and decision-making will remain under the control of the relevant human stakeholders.

Finally, OTE will contribute to the definition of the pilot validation setup, including data protection boundaries, confidentiality restrictions, performance expectations and success criteria. This will help ensure that the pilot remains feasible, safe and aligned with OTE's internal governance and regulatory obligations.

Validation and Performance Assessment

The validation of the OTE pilot will focus on assessing whether the CyberAid capabilities can support more effective detection, prioritisation and investigation of suspicious AML-related transaction patterns in a payment service provider context. The exact validation setup will depend on the final pilot approach, including the availability of synthetic, anonymised, pseudonymised or simulated data, and the internal approvals required for their use.

The performance assessment will consider both technical and operational criteria. From a technical perspective, the pilot may evaluate the ability of graph-based analytics to identify suspicious relationships between entities, the capacity of monitoring and correlation mechanisms to combine transaction-related events with technical indicators, and the usefulness of LLM-supported investigation in explaining and summarising suspicious patterns.

From an operational perspective, the pilot will assess whether the CyberAid outputs are understandable, useful and actionable for the relevant stakeholders. Particular attention will be given to the clarity of alerts, the quality of supporting evidence, the explainability of detected patterns, and the extent to which the approach can help reduce manual investigation effort. Final assessment and decision-making will remain with the appropriate human stakeholders.

Indicative validation criteria may include detection accuracy, false-positive reduction, alert prioritisation quality, investigation time reduction, explainability, auditability and suitability for AML-related reporting or documentation needs. These criteria will be refined together with the relevant OTE teams to ensure that the pilot assessment remains realistic, measurable and aligned with internal governance, privacy and regulatory requirements.

Impact

Pilot #2 is expected to demonstrate how AI-driven cybersecurity technologies can support more effective AML monitoring in payment service provider environments. By combining transaction network analysis, alert correlation and LLM-supported investigation, the pilot aims to improve the ability to detect suspicious transaction patterns that may be difficult to identify through traditional rule-based approaches alone.

The pilot can also help reduce operational effort by supporting better alert prioritisation and providing clearer context around suspicious activity. This is particularly important for AML-related investigations, where analysts need to understand not only individual transactions, but also the relationships, behaviours and technical signals behind them.

From a compliance and governance perspective, the pilot is expected to contribute to more explainable and auditable AML investigation workflows. The use of structured evidence, contextual explanations and human-in-the-loop review can support better documentation of suspicious cases, while respecting data protection, confidentiality and internal approval constraints.

Overall, the OTE pilot will provide practical feedback on how CyberAid capabilities can be applied to payment-related AML scenarios in a controlled and privacy-aware manner. The results will help refine the project's requirements for graph analytics, monitoring integration, explainability, data protection and operational validation in financial-sector environments.

5.4.3 Pilot #3 – Integrated Security Operations

Overview and Objectives

The CaixaBank pilot (Pilot 3) within the CyberAid project focuses on the validation of AI-driven risk management capabilities in a large-scale financial Security Operations Centre (SOC). The pilot is primarily centred on addressing the need for enhanced risk assessment, prioritisation, and decision support in cybersecurity operations (UC3.2), supported by advanced detection and contextualisation tasks (UC3.1).

The main objective of the pilot is to demonstrate how CyberAid technologies can support a risk-based approach to cybersecurity operations, enabling more accurate, context-aware, and consistent decision-making processes for incident handling in financial institutions.

In particular, UC3.2 aims to transform the way security events are assessed and managed by introducing AI-driven mechanisms that:

- Evaluate the risk associated with security incidents based on technical, operational, and business context.
- Support prioritisation and response decisions grounded in risk exposure.
- Provide explainable and auditable outputs, aligned with regulatory expectations.
- Enhance the efficiency and consistency of SOC workflows.

The pilot is therefore positioned as a validation environment for risk-aware SOC operations, where AI augments analyst capabilities while preserving human oversight.

Pilot Scope and Operational Context

The pilot is executed within CaixaBank's SOC environment, characterised by high volumes of security alerts, heterogeneous data sources, and strict regulatory requirements typical of the financial sector.

Within this context, the pilot integrates:

- **UC3.1**, responsible for improving the detection and correlation of security events.
- **UC3.2**, acting as the core layer for **risk analysis and decision support**.

The workflow follows a logical progression:

1. Detection and enrichment of events (UC3.1)
2. Risk assessment and prioritisation (UC3.2)

An additional step will come afterwards, for the analyst validation and response execution. This task will be undertaken by CXB, but it is out of the scope of the CYBERAID project).

This approach ensures that risk evaluation is not performed on isolated alerts but on **contextualised, enriched information**, improving the reliability of subsequent decisions.

The pilot explicitly adopts a **human-in-the-loop model**, where analysts supervise and validate AI outputs, ensuring operational control and regulatory compliance.

Use Case UC3.1 – Detection and Contextualisation

UC3.1 provides the foundational capability for the pilot by enhancing the quality and relevance of detected security events. Its role within Pilot 3 is to ensure that risk assessment in UC3.2 is based on **accurate and context-rich inputs**.

The main contributions of UC3.1 include:

- Correlation of alerts across multiple domains (network, endpoint, identity, etc.)
- Contextual enrichment using threat intelligence and asset information
- Identification of complex or multi-stage attack scenarios
- Reduction of redundant or low-value alerts

By improving signal quality, UC3.1 enables UC3.2 to operate on **meaningful and prioritised data**, reducing noise and supporting more effective risk evaluation.

Use Case UC3.2 – AI-Driven Risk Assessment and Decision Support

UC3.2 constitutes the **core of the CaixaBank pilot**, focusing on the integration of AI techniques to support risk-based decision-making in SOC operations.

The objective of this use case is to provide a structured and consistent mechanism to assess the **risk level of detected incidents**, taking into account multiple dimensions beyond technical severity. These include business impact, asset criticality, threat intelligence, and contextual indicators.

Key functionalities of UC3.2 include:

- **Risk scoring and prioritization**
Automated assessment of incidents based on a combination of technical and business factors, enabling more accurate prioritisation aligned with organisational risk exposure.
- **Context-aware analysis**
Integration of heterogeneous data sources to build a comprehensive view of each incident, supporting a deeper understanding of its potential impact.
- **Actionable recommendations**
Provision of guidance to SOC analysts regarding investigation paths and response options, linked to the assessed risk level.
- **Explainability and traceability**
Generation of transparent justifications for each risk assessment and recommendation, allowing analysts to understand and validate AI outputs.

A key aspect of UC3.2 is its focus on **consistency and standardisation**. By formalising the evaluation of risk, the system reduces variability in analyst decisions and supports more predictable and auditable SOC operations.

Integration Approach

The CyberAid components are integrated into the CaixaBank SOC ecosystem as an additional analytical layer, designed to complement existing tools without disrupting operational workflows.

The integration approach includes:

- Ingestion of alerts and contextual data from existing SOC platforms.
- Processing through CyberAid detection (UC3.1) and risk assessment (UC3.2) components.
- Generation of risk-enriched outputs and recommendations.
- Presentation of results in a format aligned with analyst workflows.

This approach ensures that CyberAid capabilities can be evaluated under realistic conditions while maintaining compatibility with enterprise environments.

Trust, Explainability and Regulatory Alignment

Given the criticality of risk management in the financial sector, the pilot places strong emphasis on **trustworthiness and regulatory compliance**, particularly within UC3.2.

The pilot addresses these aspects through:

- Transparent and explainable risk scoring mechanisms.
- Full traceability of data inputs and decision logic.
- Inclusion of confidence indicators in AI outputs.
- Mandatory human validation for critical actions.

These features support alignment with key regulatory frameworks, including:

- **NIS2 Directive** (incident handling and risk management).
- **DORA** (operational resilience and risk governance).
- **EU AI Act** (requirements for trustworthy and explainable AI systems).

Expected Outcomes

The pilot is expected to demonstrate the added value of integrating AI-driven risk assessment into SOC operations.

Key expected outcomes include:

- Improved **accuracy and consistency in incident prioritisation**.
- Reduction of **analysis time** through structured risk evaluation.
- Enhanced **decision-making confidence** for SOC analysts.
- Better alignment between **cybersecurity operations and business risk**.
- Increased preparedness for **regulatory audits and compliance requirements**.

5.4.4 Pilot #4 – Financial Market and Trading Systems

Pilot #4 focuses on the deployment and validation of CyberAid's AI-powered cybersecurity technologies within a high-frequency trading (HFT) environment. The pilot addresses the critical need to secure algorithmic trading operations, where even minor disruptions or data manipulation can lead to significant financial and systemic risks. The objective is to ensure the protection of JRC's 'proprietary trading algorithms, maintain the integrity of high-frequency market data, and safeguard the communication between trading platforms, all while preserving the ultra-low latency requirements inherent to such environments.

A key challenge to be addressed in this pilot is the integration of advanced cybersecurity monitoring without negatively impacting trading performance. In parallel, the pilot supports regulatory compliance by enabling automated incident reporting aligned with financial supervisory authorities, thereby reducing the burden of manual auditing and investigation.

A central element of this pilot is the secure and resilient data exchange between TradeStation and JRC's execution platform. This communication bridge represents the operational core of the JRC's trading lifecycle, enabling the continuous flow of market data, trading signals, execution orders, and confirmations. Given its real-time and high-frequency nature, this interface constitutes a critical attack surface. CyberAid therefore embeds advanced monitoring and anomaly detection mechanisms directly within this data flow to ensure integrity, authenticity, and timeliness without performance degradation.

Technical Contribution

The technical contribution of Pilot #4 lies in the design and implementation of a comprehensive, multi-layered cybersecurity framework that spans the entire trading lifecycle. This includes the protection of algorithm repositories, real-time monitoring of market data feeds, secure API communication, and validation of trade execution processes.

At the algorithm level, CyberAid leverages eBPF-based monitoring techniques to observe access to EasyLanguage repositories in real time. This enables the detection of unauthorized access attempts, suspicious modifications, or potential intellectual property exfiltration. By operating at the system level, this approach provides deep visibility into process behaviour while maintaining minimal performance overhead.

A key innovation of the pilot is the protection of the JRC's data exchange layer between TradeStation and execution platform. Market data feeds ingested by TradeStation are continuously analysed to detect anomalies such as irregular patterns, unexpected delays, or signs of manipulation. These data streams serve as inputs for algorithmic trading strategies, which generate execution orders transmitted through the API bridge to execution platform.

To secure this transmission layer, CyberAid performs continuous monitoring of API communications, analysing request structures, authentication behaviour, and trading patterns. This enables the identification of unauthorized API calls, abnormal order generation, or deviations from expected strategy behaviour. The system ensures that only legitimate and consistent trading instructions are transmitted to the execution platform.

In parallel, execution feedback received from the execution Platform is continuously validated. Execution confirmations, including order status, timestamps, and fill prices, are cross-checked against submitted orders to ensure consistency and integrity. Latency monitoring mechanisms measure the time between order generation and execution confirmation, allowing the detection of abnormal delays that may indicate technical failures or cyber incidents.

The framework also incorporates comprehensive telemetry and monitoring capabilities. API logs, latency metrics, and system-level events are collected and analysed in real time, enabling correlation across multiple data sources. This supports advanced forensic analysis, incident investigation, and continuous system validation, while preserving the performance requirements of high-frequency trading.

CyberAid further enhances its capabilities through the integration of AI-driven risk modeling via the OLISTIC engine. This component provides real-time threat scoring and contextual analysis of detected anomalies, enabling proactive identification of emerging risks. In addition, automated incident response mechanisms and compliance tools generate detailed audit trails and regulatory reports aligned with DORA and MiFID II requirements.

Stakeholder Contributions

The successful implementation of Pilot #4 relies on the close collaboration between the JRC's trading environment provider and CyberAid. Each partner contributes to complementary expertise and resources necessary.

The implementation of Pilot #4 will be based on close collaboration between the trading environment provider at JRC and CyberAid. Each partner contributes complementary expertise and resources required for the deployment and validation of the solution.

The trading environment provider is responsible for granting controlled access to the operational infrastructure, including trading platforms, API bridges, and market data sources. This includes both real-time and historical data, algorithm repositories, execution logs, and authentication records. The partner also provides domain expertise to assess the impact of cybersecurity mechanisms on trading performance and to validate system behaviour under realistic operating conditions.

CyberAid is responsible for the design, deployment, and integration of the cybersecurity framework. This includes the implementation of monitoring tools, anomaly detection systems, AI-based risk modelling, and automated reporting mechanisms. CyberAid also develops operational dashboards that provide real-time visibility into system performance, security status, and detected anomalies.

Validation and Performance Assessment

The validation of Pilot #4 will be based on a set of key performance indicators that capture both cybersecurity effectiveness and operational efficiency. The system is expected to achieve high accuracy in detecting anomalous or malicious events across the trading lifecycle, particularly within the data exchange layer between TradeStation and execution Platform.

A critical requirement is the preservation of ultra-low latency. The cybersecurity mechanisms must operate without introducing measurable delays that could impact trading performance. Validation therefore includes benchmarking against baseline latency metrics and continuous monitoring of execution timing.

Additional validation criteria include the accurate detection of anomalies in market data feeds, identification of abnormal execution delays, and verification of API communication integrity. The system must also demonstrate the ability to generate complete and regulatory compliant incident reports.

Validation activities are conducted through a combination of real-time monitoring, controlled simulations, and back-testing within a sandbox environment. These activities ensure that the solution meets both technical and operational requirements under realistic trading conditions.

Impact

Pilot #4 introduces a novel approach to cybersecurity in high-frequency trading environments by embedding AI-driven analytics directly into the operational data flow. A key innovation is the protection of the communication bridge between TradeStation and execution Platform, ensuring end-to-end security from market data ingestion to trade execution and confirmation.

By integrating low-level system monitoring with advanced AI-based risk modelling, the solution goes beyond traditional rule-based security approaches. It enables context-aware detection of anomalies, adaptive threat identification, and proactive risk mitigation.

The automation of incident response and regulatory reporting further bridges the gap between cybersecurity operations and compliance requirements. This enhances transparency, reduces operational overhead, and improves auditability.

The pilot demonstrates that advanced cybersecurity solutions can be effectively deployed in latency-sensitive financial environments without compromising performance. This contributes to strengthening the resilience, reliability, and trustworthiness of modern digital financial markets.

Expected Outcomes

The expected outcomes of Pilot #4 include significant improvements in the security, reliability, and transparency of algorithmic trading systems. These include enhanced protection of proprietary trading algorithms, improved detection of data anomalies and cyber threats, and increased resilience through automated failover mechanisms.

The pilot also enables comprehensive visibility into the trading lifecycle at JRC through detailed logging, monitoring, and reporting. This supports regulatory compliance while reducing the need for manual auditing and investigation.

Furthermore, the solution establishes a scalable and extensible framework that can be applied to other financial systems and domains requiring high-performance cybersecurity. By demonstrating the feasibility of integrating AI-driven security into high-frequency trading environments, Pilot #4 lays the foundation for more secure and resilient digital financial infrastructures.

5.4.5 Cross-Pilot Requirements Contributions

The pilot activities contribute to the CyberAid requirements baseline by validating the relevance of the identified requirements in different financial-sector contexts. Each pilot focuses on a specific operational environment and therefore provides input to different CyberAid capability areas. This section summarises how the four pilots contribute to the requirements catalogue and identifies the main requirement groups that are expected to be supported or validated through each pilot.

The mapping that is presented in Table 16 below is not intended to replace the detailed requirements catalogue. Instead, it provides a pilot-oriented view of the requirements and shows how the operational scenarios contribute to the refinement and validation of the CyberAid-BLUEPRINTS, CyberAid-PROACTIVE, CyberAid-MONITOR and CyberAid-REPORT areas.

Pilot	Main Contribution to Requirements	Indicative Related Requirement IDs
Pilot #1 – KM3 Client Impersonation Detection	Contributes requirements related to impersonation detection, behavioural deviation analysis, secure notification, geographic/contextual correlation, privacy-preserving federated learning, confidence scoring and quantum authentication simulation.	R-MON-15, R-MON-16, R-MON-17, R-PRO-13, R-PRO-14, R-PRO-16, R-PRO-22, R-REP-11, R-REP-17
Pilot #2 – OTE Payment Processing / iSOC	Contributes requirements related to AML-oriented transaction monitoring, graph-based modelling of financial entities and transactions, suspicious pattern detection, alert investigation, explainability, auditability and privacy-aware validation using synthetic, anonymised, pseudonymised or simulated data.	R-MON-23, R-MON-24, R-MON-25, R-MON-26, R-MON-27, R-MON-28, R-PRO-16, R-PRO-22, R-REP-07, R-REP-11, R-REP-13, R-REP-18
Pilot #3 – CXB Integrated Security Operations	Contributes requirements related to SOC alert contextualisation, behavioural analysis, multi-agent triage, AI-driven risk scoring, decision support, explainability, traceability, auditability and human-in-the-loop validation for incident prioritisation.	R-MON-06, R-MON-07, R-MON-09, R-MON-21, R-MON-30, R-MON-32, R-REP-08, R-REP-10, R-REP-11, R-REP-12, R-REP-13, R-REP-14, R-REP-17, R-REP-18, R-REP-20, R-REP-22, R-REP-23
Pilot #4 – JRC Financial Market and Trading Systems	Contributes requirements related to secure real-time data exchange, API bridge monitoring, market data integrity, execution latency monitoring, algorithm repository protection, low-latency cybersecurity, automated failover, AI risk modelling and regulatory reporting for trading environments.	R-MON-11, R-MON-12, R-MON-13, R-MON-14, R-MON-19, R-MON-20, R-MON-21, R-MON-31, R-PRO-12, R-REP-15, R-REP-16

Table 16: Cross-pilot contribution to CyberAid requirements

The pilot contributions will be further refined during the validation activities, as the technical implementation, data availability, deployment constraints and pilot scenarios become more mature. This mapping will also support later traceability work by showing which requirements are connected with concrete financial-sector validation environments.

5.5 User Stories

The user stories complement the requirements catalogue by translating selected requirements into practical scenarios from the perspective of the actors that will interact with, validate, or benefit from the CyberAid framework. While requirements describe what the platform should support, user stories help describe how these capabilities are expected to be used in realistic operational contexts.

In the scope of D2.1, user stories are used as a bridge between the requirements catalogue, the pilot scenarios, and the later validation activities. They provide a more user-oriented view of the CyberAid capabilities and help clarify the expected behaviour of the system for different roles, such as security analysts, SOC operators, compliance officers, system administrators, pilot stakeholders, and decision-makers.

The user stories were derived from the consolidated requirements catalogue and from the inputs collected from project partners and pilot stakeholders. Each user story is linked, where possible, to one or more requirements, CyberAid components, project tasks, and pilot scenarios. This mapping supports traceability and helps ensure that the requirements are not only documented, but also connected to concrete usage and validation cases.

The user stories are not intended to replace the requirements catalogue. Instead, they provide an additional layer of interpretation that can support technical design, implementation planning, pilot preparation, and validation. As the architecture and pilots become more mature, the user stories may be further refined, extended, merged, or prioritised.

5.5.1 User Stories Definition Approach

The user stories follow a common structure in order to ensure consistency across the different CyberAid capability areas. Each user story includes an identifier, title, involved role or actor, description, acceptance criteria, dependencies, and assumptions. This structure allows the consortium to capture both the expected system behaviour and the conditions under which the user story can be considered successfully fulfilled.

The acceptance criteria are particularly important because they define what needs to be achieved for the user story to be validated. These criteria will support the later definition of pilot validation activities and will help connect the requirements baseline with measurable outcomes.

User stories are derived from the identified stakeholder and pilot requirements and are intended to provide a user-centred representation of the expected CyberAid functionalities. The User Stories Matrix also maintains traceability by linking each user story to the related requirement identifiers in the “Related Req(s)” column. In addition, the matrix captures the expected business value associated with each story, highlighting the operational benefit, improvements, security enhancements, or compliance-related outcomes that CyberAid is expected to deliver. This approach helps to validate the requirements from a technical and business perspective and helps to assess how the proposed capabilities meet stakeholder needs in the different pilot environments.

5.5.2 User Stories Matrix

Table 17 presents the initial set of CyberAid user stories. The table provides a user-oriented view of the selected requirements and supports the mapping between stakeholder needs, technical capabilities and validation scenarios.

ID	User Role	User Story	Business Value	Dependencies & Assumptions	Acceptance Criteria	Related Req(s)
US-01	Security Knowledge Engineer	I want to ingest, normalize, and validate CTI sources with provenance and confidence information	So that later LLM reasoning, monitoring, and reporting outputs are based on traceable and trustworthy intelligence	Requires agreed CTI sample feeds, STIX/TAXII or equivalent formats, SKM/ontology baseline, access-control rules, and data-quality checks.	✓ At least one representative CTI feed ingested ✓ Normalized intelligence mapped to SKM entities ✓ Provenance /confidence preserved ✓ Data-quality issues logged ✓ Access-control and auditability checks passed	R-BLP-01, R-BLP-03, R-BLP-04, R-BLP-05, R-BLP-13
US-02	SOC Analyst	I want CyberAid to analyse a critical security event using multi-agent triage and contextual LLM reasoning	So that I can receive a fast, evidence-backed assessment instead of isolated alerts	Requires event queue, CRA routing, selected specialized agents, monitoring input, knowledge base/vector store, and synthetic or pilot-like critical events.	✓ Critical event injected successfully ✓ CRA classifies severity and category ✓ Relevant agents are dispatched ✓ Initial triage delivered within target latency ✓ Output includes evidence references and recommended next steps ✓ Agent trace/audit log available	R-MON-07, R-MON-09, R-MON-30, R-MON-32, R-REP-22, R-REP-23
US-03	Compliance Officer	I want CyberAid to generate DORA-aligned ICT risk metrics and regulatory reporting material from validated incidents	So that compliance reporting can be supported with structured, auditable evidence	Requires incident classification, audit logs, compliance mapping catalogue, risk metrics/KRIs, and review by legal/compliance partner.	✓ Incident mapped to relevant reporting category ✓ ICT risk metrics/KRIs generated ✓ Sample report produced ✓ Legal/compliance review completed ✓ Audit trail links report outputs to source evidence	R-REP-07, R-REP-13, R-REP-15, R-REP-19, R-REP-20
US-04	Data Protection Officer	I want to evaluate privacy-preserving federated analytics without moving sensitive financial data outside institutional boundaries	So that collaborative detection models can be validated while respecting GDPR/DPIA constraints	Requires common data/model schema, PySyft/federated workflow, privacy controls, participating institutions or simulated nodes, and agreed synthetic or	✓ Federated workflow executed across nodes ✓ No raw sensitive data leaves local environment ✓ Schema validation completed ✓ Privacy controls reviewed ✓ Data leakage test passed ✓ Model	R-PRO-13, R-PRO-16, R-PRO-22, R-BLP-06

				anonymized datasets.	performance/evaluation results documented	
US-05	Security Architect	I want to provision a digital twin test environment that replicates a selected financial service topology	So that vulnerabilities, attack paths, and response actions can be tested without impacting production	Requires reference topology, asset inventory, deployment specs, sanitized data, isolation controls, and test lifecycle orchestration.	✓ Digital twin provisioned from reference topology ✓ Service/API parity checked ✓ Isolation from production verified ✓ Traffic simulation or replay executed ✓ Test run produces traceability report ✓ Teardown and audit logs verified	R-PRO-17, R-PRO-18, R-PRO-19, R-PRO-20, R-PRO-21
US-06	Risk Manager	I want risk scores and dashboards to aggregate CTI, monitoring, pentesting, and incident evidence	So that operational and executive users can prioritize remediation actions based on consolidated risk	Requires integration of MONITOR, PROACTIVE, REPORT, OLISTIC/risk engine, dashboard/API, and agreed scoring methodology.	✓ Inputs received from at least two source modules ✓ Aggregated risk score generated ✓ Confidence score displayed where applicable ✓ Dashboard shows role-based view ✓ At least three actionable remediation items shown for high-risk case ✓ Residual-risk update recorded after treatment action	R-REP-12, R-REP-16, R-REP-17, R-REP-18, R-REP-20, R-PRO-26
US-07	Analyst / Human-in-the-Loop Reviewer	I want every AI-generated security recommendation to include explanation, evidence, confidence, and audit trail	So that I can validate, override, or approve AI-supported decisions in a controlled way	Requires LLM guardrails, evidence grounding, hallucination mitigation, dashboard/API, HITL workflow, and audit logging.	✓ Recommendation includes explanation and evidence ✓ Confidence score or uncertainty indicator shown ✓ Hallucination/grounding check applied ✓ Analyst can approve/override/reject recommendation	R-REP-10, R-REP-11, R-REP-13, R-REP-14, R-PRO-11, R-BLP-16

					✓ Decision and rationale stored in audit trail	
--	--	--	--	--	--	--

Table 17: Initial CyberAid User Stories Matrix

6. CONCLUSIONS AND NEXT STEPS

This deliverable presented the initial cybersecurity requirements analysis for the CyberAid project, focusing on the financial sector and its critical infrastructures. The analysis confirmed that financial organisations operate in an increasingly complex environment, where digital transformation, AI adoption, cloud dependencies, open banking, third-party services and evolving cyber threats create new security and resilience needs.

The requirements elicitation and analysis process provided a structured baseline for the project. The collected requirements were classified into functional and non-functional requirements and grouped according to the main CyberAid capability areas: CyberAid-BLUEPRINTS, CyberAid-PROACTIVE, CyberAid-MONITOR and CyberAid-REPORT. This structure supports readability and also helps connect the requirements with the technical work that will follow in the project.

The regulatory and standards analysis further highlighted the importance of aligning CyberAid with the operational and compliance needs of financial institutions. Frameworks such as DORA, NIS2, GDPR, PSD2, 4AMLD and the EU AI Act introduce requirements related to ICT risk management, incident reporting, data protection, third-party risk, payment security, AML monitoring, AI governance and human oversight. Relevant standards and operational frameworks provide additional guidance for implementation and validation.

The requirements catalogue produced in this deliverable will act as an initial input for several subsequent activities of the project. It will support the development of the reference architecture and blueprints, the definition of security knowledge models, the design of the LLM orchestration layer, the development of proactive, monitoring and reporting capabilities, and the integration of CyberAid components. It will also support the preparation of pilot validation activities and the mapping of technical outputs to operational needs.

The requirements baseline will be reviewed and further refined as the architecture, technical components, and pilot scenarios mature. The requirements, user stories and traceability mappings will be updated as needed with additional feedback from partners and pilot stakeholders. The user stories will be further linked to pilot scenarios and validation criteria, hence ensuring the requirements are connected to measurable outcomes and practical use cases. Future work will therefore focus on ensuring consistency between the requirements catalogue, the technical design, the implementation progress, the pilot validation and the project evaluation activities.

REFERENCES

- [1] ENISA, ENISA Threat Landscape: Finance Sector, European Union Agency for Cybersecurity, 2025. Available: <https://www.enisa.europa.eu/publications/enisa-threat-landscape-finance-sector>
- [2] ENISA, ENISA Threat Landscape 2024, European Union Agency for Cybersecurity, 2024. Available: <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2024>
- [3] EIOPA, Digital Operational Resilience Act (DORA). Available: https://www.eiopa.europa.eu/digital-operational-resilience-act-dora_en
- [4] NIST, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1, 2024. Available: <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- [5] Europol, Internet Organised Crime Threat Assessment (IOCTA) 2024, 2024. Available: <https://www.europol.europa.eu/publication-events/main-reports/internet-organised-crime-threat-assessment-iocta-2024>
- [6] OWASP, OWASP Top 10 for Large Language Model Applications, 2025. Available: <https://owasp.org/www-project-top-10-for-large-language-model-applications/>
- [7] CISA, Software Bill of Materials (SBOM). Available: <https://www.cisa.gov/topics/information-communications-technology-supply-chain-security/sbom>
- [8] European Parliament and Council of the European Union, Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA), Official Journal of the European Union, 2022. Available: <https://eur-lex.europa.eu/eli/reg/2022/2554/oj>
- [9] European Central Bank, Guide on Outsourcing Cloud Services to Cloud Service Providers, 2025. Available: https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.supervisory_guides202507.en.pdf
- [10] Bank of Greece, DORA – Digital Operational Resilience Act for the Financial Sector. Available: <https://www.bankofgreece.gr/en/main-tasks/supervision/dora-digital-operational-resilience-act-for-the-financial-sector>
- [11] BIS, Quantum-readiness for the Financial System: A Roadmap, BIS Papers No. 158, 2025. Available: <https://www.bis.org/publ/bppdf/bispap158.htm>
- [12] NIST, Transition to Post-Quantum Cryptography Standards, NIST IR 8547 Initial Public Draft, 2024. Available: <https://nvlpubs.nist.gov/nistpubs/ir/2024/NIST.IR.8547.ipd.pdf>
- [13] NIST, NIST Releases First 3 Finalized Post-Quantum Encryption Standards, 2024. Available: <https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards>
- [14] ENISA, Post-Quantum Cryptography: Integration Study, European Union Agency for Cybersecurity, 2022. Available: <https://www.enisa.europa.eu/publications/post-quantum-cryptography-integration-study>
- [15] European Commission, Recommendation on a Coordinated Implementation Roadmap for the Transition to Post-Quantum Cryptography, 2024. Available: <https://digital-strategy.ec.europa.eu/en/library/recommendation-coordinated-implementation-roadmap-transition-post-quantum-cryptography>
- [16] FS-ISAC, Navigating Cyber 2024: Annual Threat Review and Predictions, 2024. Available: <https://www.fsisac.com/navigatingcyber2024>
- [17] OASIS, STIX Version 2.1, OASIS Standard, 2021. Available: <https://www.oasis-open.org/standard/6426/>
- [18] OASIS, TAXII Version 2.1, OASIS Standard, 2021. Available: <https://www.oasis-open.org/standard/taxii-version-2-1/>
- [19] MITRE, MITRE ATT&CK®. Available: <https://attack.mitre.org/>
- [20] MITRE, Enhanced Cyber Threat Model for Financial Services Sector Institutions, 2018. Available: <https://www.mitre.org/news-insights/publication/enhanced-cyber-threat-model-financial-services-sector-institutions>
- [21] EIOPA, Guidelines on Information and Communication Technology Security and Governance, EIOPA-BoS-20/600, 2020. Available: https://www.eiopa.europa.eu/publications/guidelines-information-and-communication-technology-security-and-governance_en
- [22] European Banking Authority, Guidelines on ICT and Security Risk Management, EBA/GL/2019/04, 2019. Available: <https://www.eba.europa.eu/activities/single-rulebook/regulatory-activities/internal-governance/guidelines-ict-and-security-risk-management>
- [23] European Parliament and Council of the European Union, Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union (NIS2), Official Journal of the European Union, 2022. Available: <https://eur-lex.europa.eu/eli/dir/2022/2555/oj>
- [24] European Parliament and Council of the European Union, Regulation (EU) 2016/679, General Data Protection

- Regulation (GDPR), Official Journal of the European Union, 2016. Available: <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- [25] Ryffel, T., Trask, A., Dahl, M., Wagner, B., Mancuso, J., Rueckert, D., and Passerat-Palmbach, J., A Generic Framework for Privacy Preserving Deep Learning, arXiv:1811.04017, 2018. Available: <https://arxiv.org/abs/1811.04017>
- [26] European Parliament and Council of the European Union, Directive (EU) 2015/2366 on payment services in the internal market (PSD2), Official Journal of the European Union, 2015. Available: <https://eur-lex.europa.eu/eli/dir/2015/2366/oj>
- [27] European Commission, Commission Delegated Regulation (EU) 2018/389 supplementing PSD2 with regard to regulatory technical standards for strong customer authentication and common and secure open standards of communication, Official Journal of the European Union, 2018. Available: https://eur-lex.europa.eu/eli/reg_del/2018/389/oj
- [28] European Parliament and Council of the European Union, Directive (EU) 2015/849 on the prevention of the use of the financial system for the purposes of money laundering or terrorist financing (4AMLD), Official Journal of the European Union, 2015. Available: <https://eur-lex.europa.eu/eli/dir/2015/849/oj>
- [29] European Parliament and Council of the European Union, Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act), Official Journal of the European Union, 2024. Available: <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- [30] International Organization for Standardization, ISO/IEC 27001:2022 – Information Security, Cybersecurity and Privacy Protection – Information Security Management Systems – Requirements, ISO, 2022. Available: <https://www.iso.org/standard/82875.html>
- [31] International Organization for Standardization, ISO/IEC 27035-1:2023 – Information Technology – Information Security Incident Management – Part 1: Principles and Process, ISO, 2023. Available: <https://www.iso.org/standard/78973.html>
- [32] National Institute of Standards and Technology, The NIST Cybersecurity Framework (CSF) 2.0, NIST CSWP 29, 2024. Available: <https://doi.org/10.6028/NIST.CSWP.29>
- [33] European Central Bank, TIBER-EU Framework: How to Implement the European Framework for Threat Intelligence-Based Ethical Red Teaming, 2018. Available: <https://www.ecb.europa.eu/paym/cyber-resilience/tiber-eu/html/index.en.html>
- [34] OWASP Foundation, OWASP API Security Top 10 – 2023. Available: <https://owasp.org/API-Security/editions/2023/en/>
- [35] OWASP Foundation, OWASP Application Security Verification Standard (ASVS). Available: <https://owasp.org/www-project-application-security-verification-standard/>
- [36] Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin), DORA Overview. Available: https://www.bafin.de/DE/unternehmen-maerkte/aufsicht/alle-unternehmen/dora/ueberblick/ueberblick_node.html
- [37] Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin), Cyber Crises and Emergency Exercises under DORA. Available: https://www.bafin.de/DE/unternehmen-maerkte/aufsicht/alle-unternehmen/dora/Informationen_Cyberkrisen_Notfalluebungen/Informationen_Cyberkrisen_Notfalluebungen_node.html
- [38] Cloud Security Alliance, Anjuna – The State of Cloud and AI for Financial Services: 2026 Industry Survey Report, 2026. Available: <https://cloudsecurityalliance.org>

ANNEX I. CYBERAID REQUIREMENTS

R-BLP-01 Finance Security Knowledge Model (SKM) core entities & relations	
Req ID	R-BLP-01
Title	Finance Security Knowledge Model (SKM) core entities & relations
Description	CyberAid shall implement a finance-focused SKM capturing relationships among threats, vulnerabilities, assets, controls, and regulatory requirements, enabling consistent interpretation and analysis across the platform.
Source Type	Other
Source Reference	GA Annex 1 — T2.3 description
Validation Means	Design inspection + ontology/schema review; example instantiation tests
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	IQB
Related WP / Task	WP2 / T2.3
Notes	—

R-BLP-02 Standardized ontologies & taxonomies	
Req ID	R-BLP-02
Title	Standardized ontologies & taxonomies
Description	CyberAid shall define and publish ontologies/taxonomies that standardize security concepts across CyberAid components and support interoperability with industry schemas.
Source Type	Other
Source Reference	GA Annex 1 — T2.3 description
Validation Means	Ontology validation (consistency checks), peer review workshop, mapping report
Functional / Non-Functional	Functional

MoSCoW Priority	Must Have
Contributor	IQB
Related WP / Task	WP2 / T2.3
Notes	—

R-BLP-03 CTI Ingestion, Normalization and External Threat Intelligence Integration	
Req ID	R-BLP-03
Title	CTI Ingestion, Normalization and External Threat Intelligence Integration
Description	CyberAid shall specify and implement CTI integration mechanisms to ingest, normalize, correlate and preserve provenance/confidence for intelligence from financial ISAC feeds, STIX/TAXII sources, open-source repositories and other external threat intelligence providers.
Source Type	Other
Source Reference	GA Annex 1 — T2.3; external threat intelligence integration
Validation Means	Integration test with sample feeds; STIX/TAXII validation; data quality checks; correlation unit tests; provenance/confidence inspection.
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	IQB / UPRC
Related WP / Task	WP2 / T2.3; WP5
Notes	Consolidated from R-BLP-03, R-PRO-17

R-BLP-04 Semantic indexing for LLM reasoning	
Req ID	R-BLP-04
Title	Semantic indexing for LLM reasoning
Description	CyberAid shall support semantic indexing grounded to SKM entities to enable LLM-based reasoning over both structured CTI and unstructured sources (bulletins, reports), including extraction patterns for relevant threat knowledge.

Source Type	Other
Source Reference	GA Annex 1 — T2.3 description
Validation Means	Retrieval evaluation (RAG tests), indexed-corpus inspection, accuracy/traceability tests
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	IQB
Related WP / Task	WP2 / T2.3
Notes	—

R-BLP-05 CTI governance: integrity, access control, and auditability	
Req ID	R-BLP-05
Title	CTI governance: integrity, access control, and auditability
Description	CTI and SKM datasets shall be protected with role-based access control, integrity protection, and auditable change history (who changed what/when/why) for all model and feed configuration updates.
Source Type	Standard
Source Reference	ISO/IEC 27001:2022 (A.5/A.8/A.12 themes)
Validation Means	Access-control tests; audit-log review; tamper-evidence checks
Functional / Non-Functional	Non-Functional
MoSCoW Priority	Must Have
Contributor	IQB
Related WP / Task	WP2 / T2.3
Notes	—

R-BLP-06 Privacy-by-design for CTI enrichment & sharing	
Req ID	R-BLP-06

Title	Privacy-by-design for CTI enrichment & sharing
Description	When enriching or sharing CTI, CyberAId shall apply data minimization and privacy-by-design safeguards (e.g., pseudonymization/redaction of personal data; purpose limitation) and document confidentiality tiers for bidirectional sharing.
Source Type	Regulation
Source Reference	GDPR (Articles 5, 25, 32)
Validation Means	DPIA-aligned checklist; redaction tests; policy/controls review
Functional / Non-Functional	Non-Functional
MoSCoW Priority	Must Have
Contributor	IQB
Related WP / Task	WP2 / T2.3
Notes	—

R-BLP-07 Specialized Agent — MCP Server Deployment Pattern	
Req ID	R-BLP-07
Title	Specialized Agent — Dual-Mode Interface Contract
Description	Every specialized domain agent shall expose a standardised dual-mode interface: (a) orchestrated mode — invocable by the Central Router Agent (CRA), with Security Context enforcement and append-only audit on every invocation; and (b) standalone mode — invocable directly via CLI or REST API. MCP Server wrapping is supported as an optional external-facing adapter to expose agents to third-party consumers or standalone microservice deployment contexts where the platform runtime is not present.
Source Type	Architecture
Source Reference	T2.4 / T5.1
Validation Means	MCP interface conformance test; dual-access mode validation per agent
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have

Contributor	INNOV
Related WP / Task	T2.4, T5.1
Notes	—

R-BLP-08 Partitioned Vector Store per Agent	
Req ID	R-BLP-08
Title	Partitioned Vector Store per Agent
Description	The knowledge infrastructure shall maintain dedicated vector store partitions per agent domain (threat intel, vulnerability, compliance, behavioural, etc.) plus a shared/global partition, with all retrieval scoped by Security Context.
Source Type	Architecture
Source Reference	T2.4
Validation Means	Partition isolation test; cross-domain retrieval boundary check; Security Context scoping validation
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	INNOV
Related WP / Task	WP3-WP5
Notes	—

R-BLP-09 Concurrent Agent Operation	
Req ID	R-BLP-09
Title	Concurrent Agent Operation
Description	The system shall support concurrent request handling from both CRA and Reporting Agent to all domain agents, with stateless agent design ensuring isolation between concurrent sessions.
Source Type	Architecture
Source Reference	T2.4

Validation Means	Concurrency stress test; session isolation validation; statelessness verification
Functional / Non-Functional	Non-Functional (Performance)
MoSCoW Priority	Must Have
Contributor	INNOV
Related WP / Task	WP3-WP5
Notes	—

R-BLP-10 MCP Tool Server or CLI Interface for All External Services	
Req ID	R-BLP-10
Title	Governed Tool Adapter Layer for All External Services
Description	Domain agents shall access all Layer 2 security services (Wazuh, OLISTIC, XÉNON, DIOSCURI, CTI feeds) through a governed Tool Adapter Layer that enforces Security Context checks, JSON schema argument validation, and append-only audit logging on every external call, regardless of the underlying transport. Adapters are implemented as typed interfaces; the platform ships native typed adapters as the default. MCP Tool Server wrapping is supported as an optional integration pattern for services that expose an MCP-compatible interface or for external consumption of the adapter surface.
Source Type	Architecture
Source Reference	T2.4
Validation Means	Per-service adapter test: Security Context enforcement verified (cross-scope call rejected + audited); malformed tool arguments rejected 100% (schema validation); every successful call produces an audit record. MCP conformance test where the MCP adapter module is instantiated.
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	INNOV
Related WP / Task	WP3-WP5
Notes	—

R-BLP-11 Configuration-Driven Threat Ontology	
Req ID	R-BLP-11
Title	Configuration-Driven Threat Ontology
Description	Attack types, agent-to-attack mappings, priority weights, and thresholds shall be defined in configuration, extensible without code changes to the core pipeline.
Source Type	Architecture
Source Reference	T2.3 / T4.5
Validation Means	Adding a new attack type requires config-only changes. Peer review confirms coverage of pilot-relevant threats.
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	UPRC
Related WP / Task	WP4/WP5
Notes	—

R-BLP-12 Hybrid Retrieval (Semantic + Keyword)	
Req ID	R-BLP-12
Title	Hybrid Retrieval (Semantic + Keyword)
Description	The RAG pipeline shall support hybrid retrieval combining semantic vector similarity search with keyword-based filtering, plus re-ranking and deduplication of retrieved chunks.
Source Type	Architecture
Source Reference	T2.4
Validation Means	Retrieval quality benchmark (precision/recall); re-ranking effectiveness test; deduplication validation
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have

Contributor	INNOV
Related WP / Task	WP3-WP5
Notes	—

R-BLP-13 External Knowledge Source Integration	
Req ID	R-BLP-13
Title	External Knowledge Source Integration
Description	The knowledge infrastructure shall ingest and index content from CVE/NVD, MITRE ATT&CK, OWASP, STIX/TAXII CTI feeds, and relevant financial regulatory frameworks (DORA, GDPR, PCI DSS).
Source Type	Standard
Source Reference	MITRE ATT&CK / STIX/TAXII
Validation Means	Ingestion pipeline test per source; index completeness check; update frequency validation
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	INNOV
Related WP / Task	WP3-WP5
Notes	—

R-BLP-14 On-Premises Deployment	
Req ID	R-BLP-14
Title	On-Premises Deployment
Description	The detection framework shall be deployable entirely on-premises. No data shall leave the institution's network. Local LLM inference shall be supported as an alternative to cloud APIs.
Source Type	Pilot
Source Reference	KM3 UC#1.1

Validation Means	Docker-based deployment validated on-premises. Local LLM option tested. Pilot testing confirms no external data transmission.
Functional / Non-Functional	Non-Functional
MoSCoW Priority	Must Have
Contributor	UPRC
Related WP / Task	WP4/WP5
Notes	—

R-BLP-15 Historical Data Ingestion	
Req ID	R-BLP-15
Title	Historical Data Ingestion
Description	The system should support the ingestion of historical "suspicious activity" logs from KM3 to train the local AI model.
Source Type	Pilot
Source Reference	Data Requirement
Validation Means	Data import successful
Functional / Non-Functional	Functional
MoSCoW Priority	Could Have
Contributor	KM3
Related WP / Task	WP6/6.2
Notes	—

R-BLP-16 Data Leakage Prevention at Agent Output	
Req ID	R-BLP-16
Title	Data Leakage Prevention at Agent Output
Description	Every agent output pipeline shall include a Data Leakage Prevention check verifying that no data outside the requesting user's permitted Security

	Context scope appears in the response — mandatory for multi-tenant deployments.
Source Type	Architecture
Source Reference	T2.4
Validation Means	Cross-tenant leakage test; Security Context boundary enforcement validation
Functional / Non-Functional	Non-Functional (Security)
MoSCoW Priority	Must Have
Contributor	INNOV
Related WP / Task	WP3-WP5
Notes	—

R-BLP-17 Supply Chain and ICT Third-Party Risk Management	
Req ID	R-BLP-17
Title	Supply Chain and ICT Third-Party Risk Management
Description	The CyberAid reference architecture and blueprints shall define governance and technical requirements for identifying, classifying, monitoring, and documenting critical ICT third-party dependencies, including security requirements for external providers and supply-chain risk controls.
Source Type	Regulation / Legislation
Source Reference	NIS2 Art. 21; DORA ICT third-party risk management obligations
Validation Means	Design inspection; compliance mapping review; review of ICT third-party dependency register and security control mapping.
Functional / Non-Functional	Non-Functional
MoSCoW Priority	Should Have
Contributor	FAIR
Related WP / Task	WP2 / T2.1, T2.2
Notes	—

R-BLP-18 AI Risk Management System	
Req ID	R-BLP-18
Title	AI Risk Management System
Description	The CyberAId architecture shall include an AI risk management approach for AI-enabled components, covering lifecycle risk identification, estimation, mitigation, monitoring, and documentation of residual risks.
Source Type	Regulation / Legislation
Source Reference	AI Act Art. 9
Validation Means	Architecture review; AI risk-management checklist; review of documented risks, mitigation measures, and residual-risk records.
Functional / Non-Functional	Non-Functional
MoSCoW Priority	Should Have
Contributor	FAIR
Related WP / Task	WP2 / T2.1, T2.2, T2.4
Notes	—

R-BLP-19 Training, Validation and Test Data Governance	
Req ID	R-BLP-19
Title	Training, Validation and Test Data Governance
Description	The CyberAId architecture shall define governance requirements for training, validation, and testing datasets used by AI components, including data-quality checks, bias examination, dataset documentation, and traceability of data sources.
Source Type	Regulation / Legislation
Source Reference	AI Act Art. 10
Validation Means	Dataset governance review; sample data-quality assessment; bias/checklist review; traceability inspection.
Functional / Non-Functional	Non-Functional

MoSCoW Priority	Should Have
Contributor	FAIR
Related WP / Task	WP2 / T2.1, T2.2, T2.4
Notes	—

R-MON-01 Continuous Network Monitoring for Financial Infrastructure	
Req ID	R-MON-01
Title	Continuous Network Monitoring for Financial Infrastructure
Description	The platform shall implement continuous network monitoring to capture network traffic and application behaviour across financial infrastructure without requiring modifications to the monitored systems.
Source Type	Other
Source Reference	GA Annex 1 - T4.2 description
Validation Means	Technical review
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	K3Y
Related WP / Task	WP4, T4.2
Notes	—

R-MON-02 Network Monitoring Performance Overhead	
Req ID	R-MON-02
Title	Network Monitoring Performance Overhead
Description	The monitoring solution shall not impose performance degradation on production financial workloads.
Source Type	Other
Source Reference	GA Annex 1 - T4.2 description

Validation Means	KPI measurement
Functional / Non-Functional	Non-Functional
MoSCoW Priority	Should Have
Contributor	K3Y
Related WP / Task	WP4, T4.2
Notes	—

R-MON-03 Financial Application–Specific Monitoring Profiles	
Req ID	R-MON-03
Title	Financial Application–Specific Monitoring Profiles
Description	The platform shall provide specialized monitoring profiles for common financial applications and services, including payment processors, banking platforms, trading systems, and authentication services, ensuring visibility across the financial technology stack.
Source Type	Other
Source Reference	GA Annex 1 - T4.2 description
Validation Means	Technical review
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	K3Y
Related WP / Task	WP4, T4.2
Notes	—

R-MON-04 Real-Time Connection Tracing and Security Event Detection	
Req ID	R-MON-04
Title	Real-Time Connection Tracing and Security Event Detection
Description	The platform shall provide real-time visibility into all active network connections and their associated processes, enabling immediate detection

	of unauthorized communications and data exfiltration attempts in financial environments. The solution shall include financial-specific detection rules to monitor and detect security-relevant events, including unusual data access patterns and privilege escalation attempts.
Source Type	Other
Source Reference	GA Annex 1 - T4.2 description
Validation Means	Security testing, Functional Testing
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	K3Y
Related WP / Task	WP4, T4.2
Notes	—

R-MON-05 Network-Application Behaviour Correlation	
Req ID	R-MON-05
Title	Network-Application Behaviour Correlation
Description	The platform shall correlate network-level activities with application behaviours, creating a holistic view of system operations that enhances attack detection capabilities.
Source Type	Other
Source Reference	GA Annex 1 - T4.2 description
Validation Means	Security testing
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	K3Y
Related WP / Task	WP4, T4.2
Notes	—

R-MON-06 Behavioural Analysis Agent — UEBA	
Req ID	R-MON-06
Title	Behavioural Analysis Agent — UEBA
Description	The Behavioural Analysis Agent shall perform user and entity behaviour analytics (UEBA), detecting insider threats and anomalous behaviour patterns using deviation from established baselines.
Source Type	Architecture
Source Reference	T2.4
Validation Means	Baseline deviation detection test; insider threat scenario simulation
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	INNOV
Related WP / Task	T3.2
Notes	—

R-MON-07 Response Latency — Critical Events	
Req ID	R-MON-07
Title	Response Latency — Critical Events
Description	For CRITICAL severity events, the CRA shall deliver an initial triage assessment within 60 seconds of event receipt and a complete multi-agent synthesized finding within 5 minutes.
Source Type	Architecture
Source Reference	T2.4
Validation Means	Load test with CRITICAL event injection; latency measurement at triage and synthesis stages
Functional / Non-Functional	Non-Functional (Performance)
MoSCoW Priority	Must Have
Contributor	INNOV

Related WP / Task	WP5
Notes	—

R-MON-08 Multi-Vector Threat Coverage	
Req ID	R-MON-08
Title	Multi-Vector Threat Coverage
Description	The system shall support detection across multiple input vectors relevant to financial services, including email-based attacks, log-based anomalies, and behavioural deviations. A domain-specific threat ontology shall define attack types, detection agents, and trigger rules per vector.
Source Type	Pilot
Source Reference	KM3 UC#1.1
Validation Means	Ontology covers ≥ 20 attack types across ≥ 4 categories. KPI monitoring on synthetic and historical datasets. Pilot testing.
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	UPRC
Related WP / Task	WP4/WP5
Notes	—

R-MON-09 Real-Time Detection	
Req ID	R-MON-09
Title	Real-Time Detection
Description	Threat identification shall occur in near-real-time at the point of ingestion (email receipt, log event arrival). The detection pipeline shall meet defined latency KPIs per severity level.
Source Type	Pilot
Source Reference	KM3 UC#1.1; DORA

Validation Means	Deterministic checks < 100ms. Full pipeline with LLM < 10s (p95). Response time evaluation during development and pilot testing.
Functional / Non-Functional	Non-Functional
MoSCoW Priority	Must Have
Contributor	UPRC
Related WP / Task	WP4/WP5
Notes	—

R-MON-10 Sender / Entity Anomaly Correlation	
Req ID	R-MON-10
Title	Sender / Entity Anomaly Correlation
Description	The system shall correlate contextual metadata (e.g., geographic origin, timing, communication patterns) with known entity profiles to detect behavioural deviations.
Source Type	Pilot
Source Reference	KM3 UC#1.1
Validation Means	Anomaly detection produces accurate findings when metadata is available. Pilot validation.
Functional / Non-Functional	Functional
MoSCoW Priority	Could Have
Contributor	UPRC
Related WP / Task	WP4 / T4.5, WP6
Notes	—

R-MON-11 API Bridge Security Monitoring	
Req ID	R-MON-11
Title	API Bridge Security Monitoring

Description	Build and monitor API traffic between TradeStation and Interactive Brokers for suspicious patterns or cyber threats.
Source Type	Pilot
Source Reference	JRC Pilot 4-UC#4.1
Validation Means	API logs, anomaly detection
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	JRC
Related WP / Task	WP5 T5.1,T5.2, T5.4/WP6 T6.2,T6.5
Notes	—

R-MON-12 | Market Data Latency Monitoring

Req ID	R-MON-12
Title	Market Data Latency Monitoring
Description	Continuously monitor latency of market data feeds from IB and Tradestation.
Source Type	Pilot
Source Reference	JRC Pilot 4-UC#4.2
Validation Means	Latency benchmarks, real-time metrics
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	JRC
Related WP / Task	WP3,WP4
Notes	—

R-MON-13 | Execution Latency Monitoring

Req ID	R-MON-13
Title	Execution Latency Monitoring
Description	Monitor latency between order submission and execution confirmation from Interactive Brokers.
Source Type	Pilot
Source Reference	JRC Pilot 4-UC#4.2
Validation Means	Execution timing reports
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	JRC
Related WP / Task	WP5, T5.4
Notes	—

R-MON-14 Market Data Feed Integrity	
Req ID	R-MON-14
Title	Market Data Feed Integrity
Description	Detect anomalies or irregular patterns in TradeStation/IB data feeds.
Source Type	Pilot
Source Reference	JRC Pilot 4-UC#4.2
Validation Means	Anomaly simulation, data validation tests
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	JRC
Related WP / Task	WP4, T4.3
Notes	—

R-MON-15 Multi-Modal Impersonation Detection	
Req ID	R-MON-15
Title	Multi-Modal Impersonation Detection
Description	The system must detect suspicious behaviour in client emails/voice transcripts/ and other ways
Source Type	Pilot
Source Reference	Detect Impersonation
Validation Means	Over (>) a specific score on synthetic and historical datasets
Functional / Non-Functional	Functional
MoSCoW Priority	Could Have
Contributor	KM3
Related WP / Task	WP5/ T5.1
Notes	Reclassify. Multi-modal impersonation detection is monitoring/detection.

R-MON-16 Secure Timely Notification	
Req ID	R-MON-16
Title	Secure Timely Notification
Description	Incident alerts shall be delivered to the appropriate compliance or operational stakeholder through approved secure channels within a configurable SLA after detection.
Source Type	Pilot
Source Reference	Performance
Validation Means	Notification workflow test; channel security review; SLA measurement.
Functional / Non-Functional	Functional
MoSCoW Priority	Could Have
Contributor	KM3
Related WP / Task	WP4/T3.6

Notes	Reclassify ID Alert/notification pipeline belongs to monitoring. Replace vague 'within some time' with a defined or configurable SLA.
--------------	---

R-MON-17 Geographic Locality Analysis	
Req ID	R-MON-17
Title	Geographic Locality Analysis
Description	The system shall correlate the geographic origin of incoming communications, where lawful and technically available through metadata, with expected client context to flag unusual locality-based deviations for analyst review.
Source Type	Pilot
Source Reference	KM3 UC#1.1
Validation Means	Metadata-source verification; synthetic deviation test; alert log review.
Functional / Non-Functional	Functional
MoSCoW Priority	Could Have
Contributor	KM3
Related WP / Task	WP6 (T6.2), WP4 (T4.5)
Notes	Reclassify. Geographic locality analysis is anomaly correlation/detection.

R-MON-18 Deployment-Time Context Configuration	
Req ID	R-MON-18
Title	Deployment-Time Context Configuration
Description	The system shall accept deployment-specific context (e.g., employee directories, known domains, trusted accounts, organisational hierarchy) to improve detection accuracy. Agents shall degrade gracefully when context is absent.
Source Type	Architecture
Source Reference	T5.1
Validation Means	Schema documented. All context-dependent agents return valid results when context is missing.

Functional / Non-Functional	Functional
MoSCoW Priority	Should Have
Contributor	UPRC
Related WP / Task	WP5 / T5.1, WP6
Notes	Deployment-time context improves detection accuracy. map to monitoring/detection configuration.

R-MON-19 Algorithm Repository Monitoring	
Req ID	R-MON-19
Title	Algorithm Repository Monitoring
Description	Monitor EasyLanguage algorithm repositories for unauthorized access or modification using eBPF and AI anomaly detection.
Source Type	Pilot
Source Reference	JRC Pilot 4-UC#4.1
Validation Means	Audit logs, alert verification
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	JRC
Related WP / Task	WP4/T4.2
Notes	Algorithm repository monitoring is MONITOR/eBPF anomaly detection.

R-MON-20 Performance Impact Control	
Req ID	R-MON-20
Title	Performance Impact Control
Description	Monitoring and analysis components for trading operations shall operate within a defined performance budget, with latency and CPU overhead measured under representative workloads.

Source Type	Pilot
Source Reference	JRC Pilot 4-UC#4.1 & UC#4.2
Validation Means	Performance benchmark under representative trading workload; added latency and CPU overhead report.
Functional / Non-Functional	Non-Functional
MoSCoW Priority	Must Have
Contributor	JRC
Related WP / Task	WP4/WP5
Notes	Latency impact control is a monitoring NFR. Define measurable threshold instead of 'must not increase trading latency'.

R-MON-21 AI Risk Modelling	
Req ID	R-MON-21
Title	AI Risk Modelling
Description	Provide real-time AI-powered threat scoring for anomalies in algorithms, market data feeds, or API traffic.
Source Type	Pilot
Source Reference	JRC Pilot 4-UC#4.1 & UC#4.2
Validation Means	Risk score validation, alert correlation
Functional / Non-Functional	Functional
MoSCoW Priority	Should Have
Contributor	JRC
Related WP / Task	WP4/WP5
Notes	Reclassify. Real-time anomaly/threat scoring for market/API data is MONITOR.

R-MON-22 Common monitoring APIs and logs submission

Req ID	R-MON-22
Title	Common monitoring APIs and logs submission
Description	Given that Cyberaid may need logging across diverse systems and distributed sites, a common logging mechanism should be in place, to hold such information. A unified interface should exist in order to aid all the individual systems report and store logs from any part of the pipeline.
Source Type	Technology
Source Reference	T4.5 alert generation and notification
Validation Means	Successful integration of different types of logs
Functional / Non-Functional	Functional
MoSCoW Priority	Should Have
Contributor	HUA
Related WP / Task	WP4
Notes	—

R-MON-23 Graph-based modelling of financial entities & transactions	
Req ID	R-MON-23
Title	Graph-based modelling of financial entities & transactions
Description	CyberAid shall implement graph-based modelling of financial relationships, transactions, and entity behaviours to identify suspicious transaction networks and patterns not captured by rule-based systems.
Source Type	Other
Source Reference	GA Annex 1 — T4.1 description
Validation Means	Prototype demo on synthetic + pilot-like datasets; algorithm evaluation report
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	IQB

Related WP / Task	WP4 / T4.1
Notes	Reclassify. Graph-based modelling of financial entities & transactions is T4.1 graph analytics / fraud monitoring, so it should be MONITOR, not PROACTIVE.

R-MON-24 Neo4j + Cypher pattern matching capability	
Req ID	R-MON-24
Title	Neo4j + Cypher pattern matching capability
Description	The graph analytics engine shall use Neo4j as a core graph store and provide Cypher-based pattern matching for fraud typologies (e.g., rings, rapid fund movement, layering/structuring motifs) over temporal transaction graphs.
Source Type	Other
Source Reference	GA Annex 1 — T4.1 description
Validation Means	Query test suite; performance benchmarks; correctness checks vs ground truth
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	IQB
Related WP / Task	WP4 / T4.1
Notes	Reclassify. Neo4j + Cypher pattern matching capability is T4.1 graph analytics / fraud monitoring, so it should MONITOR, not PROACTIVE.

R-MON-25 Specialized algorithms (NetworkX/Python) for fraud detection	
Req ID	R-MON-25
Title	Specialized algorithms (NetworkX/Python) for fraud detection
Description	CyberAid shall implement specialized graph analytics algorithms (using NetworkX/Python graph libraries) for mapping transaction flows, account relationships, and temporal patterns to detect coordinated fraudulent behaviours and anomalies.
Source Type	Other

Source Reference	GA Annex 1 — T4.1 description
Validation Means	Unit tests; offline evaluation (precision/recall); replay tests
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	IQB
Related WP / Task	WP4 / T4.1
Notes	Reclassify: NetworkX/Python fraud detection algorithms is T4.1 graph analytics / fraud mo

R-MON-26 Analyst-facing graph visualization for investigations	
Req ID	R-MON-26
Title	Analyst-facing graph visualization for investigations
Description	CyberAId shall provide financial-specific graph visualizations enabling analysts to explore suspicious networks (filtering by time windows, entity types, risk scores, and path tracing) to support investigations.
Source Type	Other
Source Reference	GA Annex 1 — T4.1 description
Validation Means	UX walkthrough + analyst feedback session; functional UI tests
Functional / Non-Functional	Functional
MoSCoW Priority	Should Have
Contributor	IQB
Related WP / Task	WP4 / T4.1
Notes	Reclassify: Analyst-facing graph visualization is T4.1 graph analytics / fraud monitoring...

R-MON-27 Integration of graph intelligence with risk/LLM layers	
Req ID	R-MON-27

Title	Integration of graph intelligence with risk/LLM layers
Description	Graph-derived intelligence shall be exportable via a defined interface to (i) the risk assessment engine (OLISTIC) and (ii) the LLM orchestration layer for contextual enrichment and explanation.
Source Type	Other
Source Reference	GA Annex 1 — T4.1 description
Validation Means	API contract tests; end-to-end integration demo; traceability checks
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	IQB
Related WP / Task	WP4 / T4.1
Notes	Graph intelligence integration with risk/LLM layers is T4.1 graph analytics

R-MON-28 Real-time readiness: scalability & latency constraints	
Req ID	R-MON-28
Title	Real-time readiness: scalability & latency constraints
Description	The graph analytics pipeline shall be engineered for scale (partitioning/parallel processing as needed) to maintain near-real-time fraud pattern detection under peak transaction volumes, with defined performance KPIs and monitored SLOs.
Source Type	Regulation + General
Source Reference	DORA (ICT risk management expectations) + platform needs
Validation Means	Load/performance testing; monitoring dashboards; SLO compliance report
Functional / Non-Functional	Non-Functional
MoSCoW Priority	Must Have
Contributor	IQB
Related WP / Task	WP4 / T4.1

Notes	Real-time graph scalability/latency is T4.1 graph analytics / fraud monitorin.
--------------	--

R-MON-29 CRA — Scheduled & Periodic Analysis	
Req ID	R-MON-29
Title	CRA — Scheduled & Periodic Analysis
Description	The CRA shall support time-triggered orchestration patterns (e.g., nightly vulnerability sweeps, weekly compliance posture summaries) in addition to event-driven workflows.
Source Type	Architecture
Source Reference	T2.4
Validation Means	Scheduled job execution test; output completeness review
Functional / Non-Functional	Functional
MoSCoW Priority	Should Have
Contributor	INNOV
Related WP / Task	T2.4, T5.1
Notes	Scheduled/periodic analysis is a monitoring/detection execution mode.

R-MON-30 Threat Intelligence Agent	
Req ID	R-MON-30
Title	Threat Intelligence Agent
Description	The Threat Intelligence Agent shall perform IoC analysis, TTP mapping to MITRE ATT&CK, threat actor attribution, and campaign correlation by querying CTI feeds and the dedicated knowledge partition.
Source Type	Architecture
Source Reference	T2.4
Validation Means	IoC matching test; ATT&CK mapping accuracy evaluation; attribution scenario test

Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	INNOV
Related WP / Task	WP4
Notes	Threat Intelligence Agent supports monitoring/detection context.

R-MON-31 Automated Failover	
Req ID	R-MON-31
Title	Automated Failover
Description	Automatically switch to backup market data feeds when anomalies or latency thresholds are exceeded.
Source Type	Pilot
Source Reference	JRC Pilot 4-UC#4.2
Validation Means	Failover testing, SLA validation
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	JRC
Related WP / Task	WP4, T4.3
Notes	Automated failover supports monitoring/resilience of trading operations

R-MON-32 Multi-Agent Orchestration Pipeline	
Req ID	R-MON-32
Title	Multi-Agent Orchestration Pipeline
Description	The framework shall implement a multi-agent detection pipeline using a graph-based orchestration framework (LangGraph). Agents shall execute in parallel where possible, with results fused by an orchestrator using weighted confidence scoring.

Source Type	Architecture
Source Reference	T2.4 / T5.1
Validation Means	Pipeline compiles and executes end-to-end. State traceable through all stages. Node timings recorded.
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	UPRC
Related WP / Task	WP4/WP5
Notes	Merge of Multi-agent detection orchestration overlaps with parallel agent dispatch and monitoring pipeline logic.

R-PRO-09 Vulnerability Assessment Agent	
Req ID	R-PRO-09
Title	Vulnerability Assessment Agent
Description	The Vulnerability Assessment Agent shall perform risk scoring against CVE/NVD and MITRE databases, assess exploitability in the financial context, and generate prioritized remediation plans with rationale.
Source Type	Architecture
Source Reference	T2.4
Validation Means	CVE scoring accuracy test; remediation plan review; financial context adjustment validation
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	INNOV
Related WP / Task	WP4
Notes	—

R-PRO-10 DevSecOps & Code Analysis Agent
--

Req ID	R-PRO-10
Title	DevSecOps & Code Analysis Agent
Description	The DevSecOps Agent shall perform static and dynamic code analysis, identify security anti-patterns, and verify SDLC compliance for financial software artifacts.
Source Type	Architecture
Source Reference	T2.4 / 5.1
Validation Means	SAST/DAST integration test; anti-pattern detection benchmark; SDLC compliance checklist review
Functional / Non-Functional	Functional
MoSCoW Priority	Should Have
Contributor	INNOV
Related WP / Task	WP3
Notes	—

R-PRO-11 LLM Input Protection	
Req ID	R-PRO-11
Title	LLM Input Protection
Description	CyberAid shall protect LLM-based components against malicious or unsafe inputs by combining prompt injection detection, input validation, sanitisation, and policy-based blocking before requests are processed by domain agents.
Source Type	Architecture
Source Reference	T5.1 / LLM security guardrails
Validation Means	Prompt-injection test suite; sanitisation unit tests; policy bypass tests; evidence that blocked/sanitised inputs are logged.
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	INNOV / UPRC

Related WP / Task	WP5 / T5.1
Notes	Consolidated from R-PRO-13, R-PRO-16

R-PRO-12 API Authentication & Encryption	
Req ID	R-PRO-12
Title	API Authentication & Encryption
Description	External API integrations used by CyberAid shall enforce authenticated and encrypted communication, using approved mechanisms such as TLS and token- or certificate-based authentication.
Source Type	Pilot
Source Reference	JRC Pilot 4-UC#4.1
Validation Means	API security test; TLS configuration review; authentication/authorization test.
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	JRC
Related WP / Task	WP2,WP3,WP4
Notes	—

R-PRO-13 Privacy-Preserving Data Handling	
Req ID	R-PRO-13
Title	Privacy-Preserving Data Handling
Description	Personal data processed during email/voice/and other analysis must be encrypted and handled according to the principle of "Data Minimization."
Source Type	Legislation
Source Reference	GDPR
Validation Means	Data Protection Impact Assessment & encryption protocol verification

Functional / Non-Functional	Non-Functional
MoSCoW Priority	Must Have
Contributor	KM3
Related WP / Task	WP1/1.4,WP5/5.4, WP6/6.1
Notes	Privacy-preserving handling is a data-protection/proactive security requirement, not a blueprint.

R-PRO-14 Q2FA Protocol, API, Simulation and Pilot Integration	
Req ID	R-PRO-14
Title	Q2FA Protocol, API, Simulation and Pilot Integration
Description	CyberAid shall provide a consolidated Quantum Two-Factor Authentication (Q2FA) capability covering the protocol definition, API data model, API adaptation, simulation environment, simulation tests, and pilot application integration for the selected Q2FA use case.
Source Type	Pilot
Source Reference	UC#1.2; REQ-Q2FA-002 to REQ-Q2FA-007
Validation Means	Protocol/API review; simulation test report; API contract tests; pilot integration test; evidence that Q2FA flow completes successfully in the selected pilot context.
Functional / Non-Functional	Functional
MoSCoW Priority	Should Have
Contributor	QUBO
Related WP / Task	WP3 / WP6
Notes	Consolidated from R-BLP-21, R-BLP-22, R-BLP-23, R-PRO-14, R-PRO-15, REQ-Q2FA-007

R-PRO-15 Quantum-Resistant Cryptography Assessment and PoC	
Req ID	R-PRO-15
Title	Quantum-Resistant Cryptography Assessment and PoC

Description	CyberAid shall assess quantum-resistant cryptography options for prioritised secure connections or transactions and provide a proof-of-concept evaluation showing feasibility, integration constraints and recommended deployment approach.
Source Type	Architecture
Source Reference	T3.4
Validation Means	PoC report; cryptographic option assessment; integration feasibility review.
Functional / Non-Functional	Functional
MoSCoW Priority	Should Have
Contributor	QUBO
Related WP / Task	WP3/WP6
Notes	—

R-PRO-16 Privacy-Preserving Federated Analytics for Sensitive Data	
Req ID	R-PRO-16
Title	Privacy-Preserving Federated Analytics for Sensitive Data
Description	CyberAid shall support privacy-preserving analytics and federated training patterns for sensitive financial data, enabling model improvement or collaborative analytics without centralising raw personal or confidential data.
Source Type	Regulation
Source Reference	GDPR; GA Annex 1 privacy-preserving analytics
Validation Means	Federated learning workflow test; privacy control review; data leakage test; GDPR/DPIA checklist.
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	UPRC / UBI
Related WP / Task	WP3 / WP4
Notes	Consolidated from R-BLP-24, R-FL-02

R-PRO-17 Digital Twin Environment Provisioning	
Req ID	R-PRO-17
Title	Digital Twin Environment Provisioning
Description	CyberAId shall provide an automated mechanism to instantiate replicated digital twin environments that mirror target production service topologies, configurations, and data flows for safe, isolated testing without impacting live systems.
Source Type	Other
Source Reference	GA Annex 1 — Digital Twin / UBI Tool
Validation Means	Provisioning test on reference topology; environment diff against production spec
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	UBI
Related WP / Task	WP3 / WP5
Notes	—

R-PRO-18 Service Replication Fidelity	
Req ID	R-PRO-18
Title	Service Replication Fidelity
Description	Digital twin environments shall faithfully reproduce the behaviour of deployed services, including API interfaces, authentication flows, data schemas, and inter-service dependencies, to within a defined fidelity tolerance, enabling representative security and functional testing.
Source Type	Other
Source Reference	GA Annex 1 — Digital Twin / UBI Tool
Validation Means	Fidelity benchmarking report; API parity tests; service contract validation
Functional / Non-Functional	Functional

MoSCoW Priority	Must Have
Contributor	UBI
Related WP / Task	WP3 / WP5
Notes	—

R-PRO-19 Network Traffic Simulation in Twin	
Req ID	R-PRO-19
Title	Network Traffic Simulation in Twin
Description	The digital twin platform shall support configurable network traffic simulation, including synthetic load generation, latency injection, packet loss, and replay of captured production traffic traces, to enable realistic security and performance testing scenarios.
Source Type	Other
Source Reference	GA Annex 1 — Digital Twin / UBI Tool
Validation Means	Traffic replay test; load generation verification; latency/packet-loss injection demo
Functional / Non-Functional	Functional
MoSCoW Priority	Should Have
Contributor	UBI
Related WP / Task	WP3 / WP5
Notes	—

R-PRO-20 Test Lifecycle Orchestration	
Req ID	R-PRO-20
Title	Test Lifecycle Orchestration
Description	CyberAId shall provide an orchestration layer for the full test lifecycle within digital twin environments, covering test plan execution, result capture, environment teardown, and auditability of test runs, with traceability from test case to requirement.

Source Type	Standard
Source Reference	ISO/IEC 29119 (Software Testing)
Validation Means	End-to-end test run with traceability report; teardown verification; audit log review
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	UBI
Related WP / Task	WP3 / WP5
Notes	—

R-PRO-21 Digital Twin Isolation & Security	
Req ID	R-PRO-21
Title	Digital Twin Isolation & Security
Description	Digital twin environments shall be strictly isolated from production systems and networks, with enforced access controls, secrets management for any injected credentials, and automated sanitisation of any production data used to seed test environments, ensuring no risk of data leakage or cross-environment interference.
Source Type	Regulation
Source Reference	DORA Article 25 — Digital Operational Resilience Testing
Validation Means	Isolation penetration test; access-control audit; data-sanitisation verification
Functional / Non-Functional	Non-Functional
MoSCoW Priority	Must Have
Contributor	UBI
Related WP / Task	WP3 / WP5
Notes	—

R-PRO-22 Common data and model structure	
Req ID	R-PRO-22
Title	Common data and model structure
Description	In the case of federated learning across different entities, in order to jointly train AI models on a common task, the data formats and selected inputs for the given model need to agree
Source Type	Technology
Source Reference	T3.5
Validation Means	Predeployment checks, schema validators and agreement on a schema
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	HUA
Related WP / Task	WP3 and WP6
Notes	—

R-PRO-23 Automated Attack Surface Discovery	
Req ID	R-PRO-23
Title	Automated Attack Surface Discovery
Description	The pentesting module shall automatically enumerate the attack surface of CyberAId-monitored financial services, including exposed APIs, network services, authentication endpoints, and third-party integrations, producing a continuously updated attack surface inventory for each pentest cycle.
Source Type	Regulation
Source Reference	DORA Article 26 — Threat-Led Penetration Testing
Validation Means	Coverage comparison against known asset inventory; asset discovery accuracy metrics
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have

Contributor	UBI
Related WP / Task	WP4 / WP5
Notes	—

R-PRO-24 AI-Driven Exploit Chain Generation	
Req ID	R-PRO-24
Title	AI-Driven Exploit Chain Generation
Description	The pentesting module shall leverage AI reasoning (hex-clad-ai) to generate multi-step exploit chains by chaining discovered vulnerabilities, misconfigurations, and credential weaknesses into realistic attack scenarios representative of advanced persistent threats targeting financial infrastructure.
Source Type	Other
Source Reference	GA Annex 1 — Pentesting / hex-clad-ai Tool
Validation Means	Exploit chain demo against digital twin; chain-generation accuracy review by security expert
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	UBI
Related WP / Task	WP4 / WP5
Notes	—

R-PRO-25 Continuous & Scheduled Pentest Execution	
Req ID	R-PRO-25
Title	Continuous & Scheduled Pentest Execution
Description	The pentesting module shall support both on-demand and scheduled pentest execution against digital twin environments, with configurable scope, intensity, and technique profiles, and shall integrate with the CyberAid CI/CD pipeline to trigger security regression tests on each deployment.

Source Type	Other
Source Reference	GA Annex 1 — Pentesting / hex-clad-ai Tool
Validation Means	Scheduled run logs; CI/CD integration demo; scope/intensity configuration verification
Functional / Non-Functional	Functional
MoSCoW Priority	Should Have
Contributor	UBI
Related WP / Task	WP4 / WP5
Notes	—

R-PRO-26 Pentest Finding Normalisation & Risk Integration	
Req ID	R-PRO-26
Title	Pentest Finding Normalisation & Risk Integration
Description	Pentest findings shall be normalised to a standard vulnerability taxonomy (CVE/CWE/CVSS) and automatically ingested into the CyberAid risk management engine (OLISTIC) to enrich the holistic risk score and trigger remediation workflows, with bidirectional traceability between findings and risk records.
Source Type	Standard
Source Reference	PTES / OWASP Testing Guide v4.2
Validation Means	Taxonomy mapping accuracy test; OLISTIC ingestion end-to-end demo; traceability audit
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	UBI
Related WP / Task	WP4 / WP5
Notes	—

R-PRO-27 Regulatory Compliance Mapping of Pentest Findings	
Req ID	R-PRO-27
Title	Regulatory Compliance Mapping of Pentest Findings
Description	Add a cross-reference that vulnerability findings generated by proactive testing/pentesting should provide evidence usable for CRA vulnerability handling, coordinated disclosure, SBOM maintenance, and remediation tracking through R-REP-20.
Source Type	Regulation
Source Reference	NIS2 Article 21 — Security Measures; DORA Article 26
Validation Means	Mapping accuracy review vs regulatory control catalogue; compliance report validation
Functional / Non-Functional	Functional
MoSCoW Priority	Should Have
Contributor	UBI
Related WP / Task	WP4 / WP5
Notes	—

R-REP-01 Forensic Evidence Collection with Chain of Custody	
Req ID	R-REP-01
Title	Forensic Evidence Collection with Chain of Custody
Description	The system will implement evidence collection mechanisms designed to preserve volatile data from financial systems while maintaining chain of custody and forensic soundness.
Source Type	Other
Source Reference	GA Annex 1 - T5.4 description
Validation Means	Forensic review
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have

Contributor	K3Y
Related WP / Task	WP5, T5.4
Notes	—

R-REP-02 Financial-Specific Forensic Workflows and Timeline Reconstruction	
Req ID	R-REP-02
Title	Financial-Specific Forensic Workflows and Timeline Reconstruction
Description	The platform shall provide forensic analysis workflows for reconstructing security incidents in financial environments, combining multiple data sources into coherent incident narratives.
Source Type	Other
Source Reference	GA Annex 1 - T5.4 description
Validation Means	Functional Testing
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	K3Y
Related WP / Task	WP5, T5.4
Notes	—

R-REP-03 Threat Intelligence Integration for Forensic Contextualization	
Req ID	R-REP-03
Title	Threat Intelligence Integration for Forensic Contextualization
Description	The system shall integrate with threat intelligence repositories in order to contextualize forensic findings within the broader financial threat landscape, identifying potential threat actors and their tactics, techniques, and procedures.
Source Type	Other
Source Reference	GA Annex 1 - T5.4 description

Validation Means	Functional Testing
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	K3Y
Related WP / Task	WP5, T5.4
Notes	—

R-REP-04 Financial Specific Root Cause Analysis	
Req ID	R-REP-04
Title	Financial Specific Root Cause Analysis
Description	The platform shall support root cause analysis specifically adapted for financial environments, identifying underlying vulnerabilities and recommending remediation actions.
Source Type	Other
Source Reference	GA Annex 1 - T5.4 description
Validation Means	Technical review
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	K3Y
Related WP / Task	WP5, T5.4
Notes	—

R-REP-05 Privacy-Preserving Forensic Investigation	
Req ID	R-REP-05
Title	Privacy-Preserving Forensic Investigation
Description	The forensic toolkit shall implement privacy-preserving investigation techniques while maintaining GDPR compliance.

Source Type	Legislation
Source Reference	GDPR
Validation Means	Regulatory compliance testing
Functional / Non-Functional	Non-Functional
MoSCoW Priority	Must Have
Contributor	K3Y
Related WP / Task	WP5, T5.4
Notes	—

R-REP-06 Result Aggregation, Corroboration and Conflict Resolution	
Req ID	R-REP-06
Title	Result Aggregation, Corroboration and Conflict Resolution
Description	The CyberAId orchestration layer shall aggregate multi-agent outputs, corroborate findings across sources, resolve conflicts, and produce a unified explanation and prioritised result for downstream reporting and response.
Source Type	Architecture
Source Reference	T5.1 / LLM orchestration
Validation Means	Multi-agent aggregation test; conflict-resolution test cases; evidence comparison; explanation review.
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	INNOV / UPRC
Related WP / Task	WP5 / T5.1
Notes	Consolidated from R-PRO-18, R-REP-06

R-REP-07 Compliance Verification Agent	
Req ID	R-REP-07

Title	Compliance Verification Agent
Description	The Compliance Verification Agent shall map identified security events and posture assessments against DORA, GDPR, and PCI DSS regulatory frameworks, flagging gaps and generating compliance evidence artifacts.
Source Type	Legislation
Source Reference	DORA; GDPR
Validation Means	Framework mapping accuracy test; gap detection scenario validation; artifact completeness review
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	INNOV
Related WP / Task	WP1/WP3
Notes	—

R-REP-08 Incident Response Agent	
Req ID	R-REP-08
Title	Incident Response Agent
Description	The Incident Response Agent shall generate containment strategies, recommend playbook activation, and support recovery planning for confirmed incidents.
Source Type	Architecture
Source Reference	T2.4
Validation Means	Containment strategy review; playbook recommendation accuracy test
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	INNOV
Related WP / Task	T5.4
Notes	—

R-REP-09 Forensic Analysis Agent	
Req ID	R-REP-09
Title	Forensic Analysis Agent
Description	The Forensic Analysis Agent shall perform timeline reconstruction, support evidence preservation procedures, and produce forensic narratives meeting evidentiary standards for financial-sector regulatory obligations.
Source Type	Architecture
Source Reference	T2.4
Validation Means	Timeline reconstruction accuracy test; evidentiary standard review
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	INNOV
Related WP / Task	T5.4
Notes	—

R-REP-10 Hallucination Mitigation	
Req ID	R-REP-10
Title	Hallucination Mitigation
Description	The output processing pipeline shall include hallucination mitigation mechanisms (e.g., evidence grounding checks, source attribution validation) before delivering findings to orchestrators or users.
Source Type	Architecture
Source Reference	T2.4
Validation Means	Grounding check accuracy test; attribution validation on sample findings
Functional / Non-Functional	Non-Functional (Security)
MoSCoW Priority	Must Have

Contributor	INNOV
Related WP / Task	WP3-WP5
Notes	—

R-REP-11 Explainable AI Output for Security and Identity Decisions	
Req ID	R-REP-11
Title	Explainable AI Output for Security and Identity Decisions
Description	CyberAid shall provide explainable AI outputs for security, fraud and identity-verification decisions, including the evidence used, confidence score or rationale, source citations where available, and auditability suitable for regulatory and operational review.
Source Type	Legislation
Source Reference	MiFID II; GDPR; pilot identity-verification needs
Validation Means	Explanation completeness test; user review with analysts; audit trail check; controlled test cases for evidence traceability.
Functional / Non-Functional	Non-Functional
MoSCoW Priority	Must Have
Contributor	UPRC / KM3
Related WP / Task	WP5 / T5.1 / T5.3 / T5.4
Notes	Consolidated from R-BLP-07, R-REP-12

R-REP-12 Actionable Response Recommendations	
Req ID	R-REP-12
Title	Actionable Response Recommendations
Description	Detection outputs shall be translated into prioritised, role-appropriate actionable tasks for end users (e.g., block, quarantine, escalate, investigate).
Source Type	Pilot
Source Reference	KM3 UC#1.1; All Pilots

Validation Means	Pilot testing with operational staff; minimum 3 actionable items per high-severity alert
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	UPRC
Related WP / Task	WP4/WP5
Notes	—

R-REP-13 Full Audit Trail and AI System Record-Keeping	
Req ID	R-REP-13
Title	Full Audit Trail and AI System Record-Keeping
Description	Extend audit-trail/logging coverage to include AI Act Art. 12 record-keeping, including input references, verification timestamps, output traceability, user/system actions, and legally required retention periods.
Source Type	Legislation
Source Reference	DORA; MiFID II; pilot compliance needs
Validation Means	Audit log completeness test; log integrity verification; sample incident reconstruction; SIEM integration check where applicable.
Functional / Non-Functional	Non-Functional
MoSCoW Priority	Must Have
Contributor	UPRC / KM3 / JRC
Related WP / Task	WP5 / T5.1 / T5.4 / WP6
Notes	Consolidated from R-BLP-12, R-REP-14, R-REP-16

R-REP-14 REST API and Analyst Dashboard	
Req ID	R-REP-14
Title	REST API and Analyst Dashboard

Description	The framework shall expose a REST API for data submission and provide a web dashboard for analysts to inspect findings, filter by severity, and review explanations.
Source Type	Architecture
Source Reference	T5.1 / T6.2
Validation Means	API functional tests. UI usability testing with pilot staff.
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	UPRC
Related WP / Task	WP5 / T5.1, WP6 / T6.2
Notes	—

R-REP-15 Regulatory Incident and Breach Reporting	
Req ID	R-REP-15
Title	Regulatory Incident and Breach Reporting
Description	Extend regulatory reporting to include NIS2 Art. 23 incident notification timelines and GDPR Arts. 33–34 personal data breach notification requirements, including early warning, full notification, final report, and affected-data-subject communication where applicable.
Source Type	Pilot
Source Reference	NIS2 Art. 23; GDPR Arts. 33–34
Validation Means	Regulatory report validation
Functional / Non-Functional	Non-Functional
MoSCoW Priority	Should Have
Contributor	JRC
Related WP / Task	WP4/WP5
Notes	—

R-REP-16 Executive, KPI and Risk Dashboard	
Req ID	R-REP-16
Title	Executive, KPI and Risk Dashboard
Description	CyberAId shall provide role-based dashboards and reports summarising security incidents, activity volumes, false-positive rates, algorithm security status, API health, market/execution latency, alerts, risk posture and remediation status for technical and executive users.
Source Type	Pilot
Source Reference	KM3 UC#1.1; JRC UC#4.1-4.2; risk dashboard needs
Validation Means	Dashboard demonstration; role-based view test; KPI verification; UX walkthrough with pilot users.
Functional / Non-Functional	Functional
MoSCoW Priority	Should Have
Contributor	KM3 / JRC / UBI
Related WP / Task	WP4 / WP5 / WP6
Notes	Consolidated from R-BLP-14, R-REP-18, R-RM-03

R-REP-17 Risk-Weighted Confidence Scoring	
Req ID	R-REP-17
Title	Risk-Weighted Confidence Scoring
Description	The system must provide a "Confidence Score" (0-100%) for identity verification and/or other fraud detection likelihood, allowing staff to quickly assess risk levels.
Source Type	Pilot
Source Reference	Usability
Validation Means	UI Usability Testing & Backend Accuracy Validation
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have

Contributor	KM3
Related WP / Task	WP5/5.2,5.4
Notes	Reclassify to Report

R-REP-18 Holistic Multi-Source Risk Aggregation	
Req ID	R-REP-18
Title	Holistic Multi-Source Risk Aggregation
Description	The OLISTIC risk management engine shall aggregate risk signals from all CyberAid modules — CTI feeds, graph analytics, LLM detections, SIEM alerts, and pentest findings — into a unified, continuously updated risk posture model, applying a configurable weighting and aggregation methodology documented in the platform design.
Source Type	Other
Source Reference	GA Annex 1 — Risk Management / OLISTIC Tool
Validation Means	Integration test with all source modules; aggregated risk score verification; methodology review
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	UBI
Related WP / Task	WP4
Notes	—

R-REP-19 DORA/NIS2 ICT Risk Management Metrics and Reporting	
Req ID	R-REP-19
Title	DORA/NIS2 ICT Risk Management Metrics and Reporting
Description	Extend the requirement to cover both DORA ICT risk-management framework documentation and NIS2 Art. 21 ICT risk-management measures, including risk analysis, incident handling, business continuity, supply-chain security, and traceability of reporting outputs to regulatory requirements.

Source Type	Regulation
Source Reference	DORA Art. 6; NIS2 Art. 21
Validation Means	DORA compliance checklist review; KRI configuration test; sample report review by legal/compliance
Functional / Non-Functional	Non-Functional
MoSCoW Priority	Must Have
Contributor	UBI
Related WP / Task	WP4
Notes	—

R-REP-20 Remediation Tracking and Vulnerability Disclosure Evidence	
Req ID	R-REP-20
Title	Remediation Tracking and Vulnerability Disclosure Evidence
Description	Extend remediation workflow to cover CRA vulnerability handling, coordinated disclosure procedures, SBOM maintenance where applicable, vulnerability remediation evidence, and linkage between technical findings and regulatory follow-up actions.
Source Type	Standard
Source Reference	ISO/IEC 27005:2022 — Information Security Risk Management
Validation Means	Workflow E2E test; audit trail completeness check; residual-risk recalculation verification
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	UBI
Related WP / Task	WP4
Notes	—

R-REP-21 | Risk Trend Analysis & Predictive Risk Indicators

Req ID	R-REP-21
Title	Risk Trend Analysis & Predictive Risk Indicators
Description	OLISTIC shall analyse historical risk data to compute risk trend indicators and, where feasible, provide predictive risk signals using ML-based forecasting, alerting operators to deteriorating risk postures before threshold breaches occur, with explainable outputs traceable to contributing risk factors.
Source Type	Regulation
Source Reference	NIS2 Article 21; AI Act Article 9 — Risk Management
Validation Means	Trend accuracy validation on historical pilot data; explainability review; alert threshold test
Functional / Non-Functional	Functional
MoSCoW Priority	Could Have
Contributor	UBI
Related WP / Task	WP4
Notes	—

R-REP-22 CRA — Event Triage & Classification	
Req ID	R-REP-22
Title	CRA — Event Triage & Classification
Description	The CRA shall receive security events from the message queue, perform initial severity assessment and incident categorization, and estimate potential impact before dispatching to specialized agents.
Source Type	Architecture
Source Reference	T2.4 / T5.1
Validation Means	Design inspection + agent routing tests; end-to-end triage scenario validation
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have

Contributor	INNOV
Related WP / Task	T2.4, T5.1
Notes	Reclassify to Report

R-REP-23 CRA — Multi-Agent Parallel Dispatch	
Req ID	R-REP-23
Title	CRA — Multi-Agent Parallel Dispatch
Description	The CRA shall be capable of simultaneously activating multiple specialized agents when an incident requires multi-domain analysis, coordinating parallel execution and collecting partial results as they arrive.
Source Type	Architecture
Source Reference	T2.4
Validation Means	Concurrent dispatch test; multi-agent integration test with timing validation
Functional / Non-Functional	Functional
MoSCoW Priority	Must Have
Contributor	INNOV
Related WP / Task	T2.4, T5.1
Notes	Reclassify to Report

R-REP-24 DPIA Support	
Req ID	R-REP-24
Title	DPIA Support
Description	The CyberAid reporting capability shall support the production of structured documentation artifacts for Data Protection Impact Assessments, including processing descriptions, necessity and proportionality considerations, risk identification, and mitigation measures.
Source Type	Regulation / Legislation
Source Reference	GDPR Art. 35

Validation Means	Compliance artifact review; DPIA template/output inspection; legal partner validation.
Functional / Non-Functional	Functional
MoSCoW Priority	Could Have
Contributor	FAIR
Related WP / Task	WP2 / T2.1; WP5 / REPORT
Notes	—

ANNEX II. CYBERAID USER STORIES

US-BLP-01 CTI Ingestion, Normalization and Knowledge Grounding	
US-ID	US-BLP-01
Title	CTI Ingestion, Normalization and Knowledge Grounding
Role(s)	Security Knowledge Engineer
Description	As a Security Knowledge Engineer, I want to ingest, normalize, and validate CTI sources with provenance and confidence information, so that later LLM reasoning, monitoring, and reporting outputs are based on traceable and trustworthy intelligence.
Acceptance Criteria	✓ At least one representative CTI feed ingested ✓ Normalized intelligence mapped to SKM entities ✓ Provenance / confidence preserved ✓ Data-quality issues logged ✓ Access-control and auditability checks passed
Dependencies	Requires agreed CTI sample feeds, STIX/TAXII or equivalent formats, SKM/ontology baseline, access-control rules, and data-quality checks.
Assumptions	CTI feeds are available in STIX/TAXII format or equivalent. SKM ontology baseline is defined and accessible. Access-control rules have been agreed by the consortium.
Related Requirements	R-BLP-01, R-BLP-03, R-BLP-04, R-BLP-05, R-BLP-13

US-MON-01 Multi-Agent Security Event Triage	
US-ID	US-MON-01

Title	Multi-Agent Security Event Triage
Role(s)	SOC Analyst
Description	As a SOC Analyst, I want CyberAId to analyse a critical security event using multi-agent triage and contextual LLM reasoning, so that I can receive a fast, evidence-backed assessment instead of isolated alerts.
Acceptance Criteria	✓ Critical event injected successfully ✓ CRA classifies severity and category ✓ Relevant agents are dispatched ✓ Initial triage delivered within target latency ✓ Output includes evidence references and recommended next steps ✓ Agent trace/audit log available
Dependencies	Requires event queue, CRA routing, selected specialized agents, monitoring input, knowledge base/vector store, and synthetic or pilot-like critical events.
Assumptions	The CRA and at least two specialized agents are deployed and operational. A synthetic or pilot-like critical event dataset is available for testing. The knowledge base and vector store contain relevant indexed content.
Related Requirements	R-MON-07, R-MON-09, R-MON-30, R-MON-32, R-REP-22, R-REP-23

US-REP-01 DORA-Aligned Regulatory Reporting from Incidents	
US-ID	US-REP-01
Title	DORA-Aligned Regulatory Reporting from Incidents
Role(s)	Compliance Officer
Description	As a Compliance Officer, I want CyberAId to generate DORA-aligned ICT risk metrics and regulatory reporting material from validated incidents, so that compliance reporting can be supported with structured, auditable evidence.
Acceptance Criteria	✓ Incident mapped to relevant reporting category ✓ ICT risk metrics/KRIs generated ✓ Sample report produced ✓ Legal/compliance review completed ✓ Audit trail links report outputs to source evidence
Dependencies	Requires incident classification, audit logs, compliance mapping catalogue, risk metrics/KRIs, and review by legal/compliance partner.

Assumptions	Incident classification and audit logging are operational. DORA compliance mapping catalogue is defined. Legal/compliance partner is available for review. Reporting output templates are agreed within the consortium.
Related Requirements	R-REP-07, R-REP-13, R-REP-15, R-REP-19, R-REP-20

US-PRO-01 Privacy-Preserving Federated Analytics Validation	
US-ID	US-PRO-01
Title	Privacy-Preserving Federated Analytics Validation
Role(s)	Data Protection Officer
Description	As a Data Protection Officer, I want to evaluate privacy-preserving federated analytics without moving sensitive financial data outside institutional boundaries, so that collaborative detection models can be validated while respecting GDPR/DPIA constraints.
Acceptance Criteria	✓ Federated workflow executed across nodes ✓ No raw sensitive data leaves local environment ✓ Schema validation completed ✓ Privacy controls reviewed ✓ Data leakage test passed ✓ Model performance/evaluation results documented
Dependencies	Requires common data/model schema, PySyft/federated workflow, privacy controls, participating institutions or simulated nodes, and agreed synthetic or anonymized datasets.
Assumptions	At least two participating nodes (or simulated nodes) are available. Synthetic or anonymized datasets are agreed and prepared. PySyft or equivalent federated learning framework is deployed. GDPR/DPIA requirements are documented.
Related Requirements	R-PRO-13, R-PRO-16, R-PRO-22, R-BLP-06

US-PRO-02 Digital Twin Environment Provisioning for Security Testing	
US-ID	US-PRO-02
Title	Digital Twin Environment Provisioning for Security Testing
Role(s)	Security Architect

Description	As a Security Architect, I want to provision a digital twin test environment that replicates a selected financial service topology, so that vulnerabilities, attack paths, and response actions can be tested without impacting production.
Acceptance Criteria	✓ Digital twin provisioned from reference topology ✓ Service/API parity checked ✓ Isolation from production verified ✓ Traffic simulation or replay executed ✓ Test run produces traceability report ✓ Teardown and audit logs verified
Dependencies	Requires reference topology, asset inventory, deployment specs, sanitized data, isolation controls, and test lifecycle orchestration.
Assumptions	A reference topology and asset inventory for the target pilot environment are available. Sanitized or synthetic data is agreed for seeding the twin. Isolation mechanisms are in place and enforceable. Test lifecycle orchestration tooling is operational.
Related Requirements	R-PRO-17, R-PRO-18, R-PRO-19, R-PRO-20, R-PRO-21

US-REP-02 Consolidated Risk Scoring and Remediation Dashboard	
US-ID	US-REP-02
Title	Consolidated Risk Scoring and Remediation Dashboard
Role(s)	Risk Manager
Description	As a Risk Manager, I want risk scores and dashboards to aggregate CTI, monitoring, pentesting, and incident evidence, so that operational and executive users can prioritize remediation actions based on consolidated risk.
Acceptance Criteria	✓ Inputs received from at least two source modules ✓ Aggregated risk score generated ✓ Confidence score displayed where applicable ✓ Dashboard shows role-based view ✓ At least three actionable remediation items shown for high-risk case ✓ Residual-risk update recorded after treatment action
Dependencies	Requires integration of MONITOR, PROACTIVE, REPORT, OLISTIC/risk engine, dashboard/API, and agreed scoring methodology.
Assumptions	At least two CyberAid modules (e.g., MONITOR and PROACTIVE/pentesting) are providing risk signals. The OLISTIC risk

	engine and dashboard/API are operational. An agreed weighting and aggregation methodology is documented.
Related Requirements	R-REP-12, R-REP-16, R-REP-17, R-REP-18, R-REP-20, R-PRO-26

US-REP-03 Explainable AI Recommendations with Human-in-the-Loop Control	
US-ID	US-REP-03
Title	Explainable AI Recommendations with Human-in-the-Loop Control
Role(s)	Analyst / Human-in-the-Loop Reviewer
Description	As an Analyst/Human-in-the-Loop Reviewer, I want every AI-generated security recommendation to include explanation, evidence, confidence, and audit trail, so that I can validate, override, or approve AI-supported decisions in a controlled way.
Acceptance Criteria	✓ Recommendation includes explanation and evidence ✓ Confidence score or uncertainty indicator shown ✓ Hallucination/grounding check applied ✓ Analyst can approve/override/reject recommendation ✓ Decision and rationale stored in audit trail
Dependencies	Requires LLM guardrails, evidence grounding, hallucination mitigation, dashboard/API, HITL workflow, and audit logging.
Assumptions	LLM orchestration layer and guardrails are operational. Evidence grounding and hallucination mitigation mechanisms are implemented. The dashboard/API exposes HITL approve/override/reject controls. Audit logging is in place and tested.
Related Requirements	R-REP-10, R-REP-11, R-REP-13, R-REP-14, R-PRO-11, R-BLP-16

