



AI-DRIVEN CYBERSECURITY FOR FINANCIAL SERVICE PROVIDERS

D2.4 – SECURITY KNOWLEDGE MODELS AND CTI INTEGRATION FRAMEWORK

Lead Beneficiary	INQBIT INNOVATIONS SRL (IQB)
Work Package Ref.	WP2-LLM-Powered Reference Architecture, Specification and Systems Integration
Task Ref.	T2.3-Security Knowledge Modelling and CTI Feeds Specification
Deliverable Title	Security Knowledge Models and CTI Integration Framework
Due Date	30/09/2026
Delivered Date	30/09/2026
Revision Number	V1.0
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 ECCC 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
IQB	Lead Beneficiary	Cosmina Stalidi
K3Y	Contributor	Panagiotis Radoglou Grammatikis, Evangelia Pantraki, Christos Dalamangas, Nikolaos Ntampakis, Dimitrios Balompas, Antonios Sarigiannidis
INNOV	Internal Reviewer	George Fatouros
UBI	Contributor	Nikos Papageorgopoulos
UBI, GFT	Quality Assurance	Louiza Kachrimani, Michele Canteri

REVISION HISTORY

Version	Date	Partner(s)	Description
v0.1	15/04/26	IQB	1 st draft version of ToC
v0.2	07/07/26	IQB	Final version of ToC
v0.3	16/07/26	IQB, K3Y	Initial input in Sections 2 and 3
v0.4	10/08/26	IQB, K3Y	Initial input in Sections 4-5
v0.5	21/08/26	IQB, UBI	Resolved comments by UBI
v0.6	07/09/26	IQB, UPRC	Resolved comments by UPRC
v0.7	14/09/26	IQB, INNOV	Resolved comments by INNOV
v0.8	15/09/26	IQB	1 st draft version for internal review
v0.9	24/09/26	IQB	Draft version for QA
v1.0	30/09/26	UBI	Final quality check and submission to the EU

¹ Lead Beneficiary, Contributor, Internal Reviewer, Quality Assurance

² Can be left void

LIST OF ABBREVIATIONS

Acronym	Definition
CAPEC	Common Attack Pattern Enumeration and Classification
CERT	Computer Emergency Response Team
CRA	Cyber Resilience Act
CTI	Cyber Threat Intelligence
CVE	Common Vulnerabilities and Exposures
CVSS	Common Vulnerability Scoring System
CWE	Common Weakness Enumeration
DORA	Digital Operational Resilience Act
EMTD	Email Threat Detection
FS-ISAC	Financial Services Information Sharing and Analysis Center
GDPR	General Data Protection Regulation
ICT	Information and Communication Technology
IoC	Indicator of Compromise
ISAC	Information Sharing and Analysis Center
ISO/IEC	International Organization for Standardization / International Electrotechnical Commission
LLM	Large Language Model
MITRE ATT&CK	MITRE Adversarial Tactics, Techniques, and Common Knowledge
NIS2	Network and Information Systems Directive 2
NIST	National Institute of Standards and Technology
NVD	National Vulnerability Database
OWASP	Open Worldwide Application Security Project

Acronym	Definition
RAG	Retrieval-Augmented Generation
SKM	Security Knowledge Model
SOAR	Security Orchestration, Automation and Response
SOC	Security Operations Centre
STIX	Structured Threat Information Expression
TAXII	Trusted Automated Exchange of Intelligence Information
TLP	Traffic Light Protocol
UUID	Universally Unique Identifier

EXECUTIVE SUMMARY

The increasing complexity of cyber threats targeting the financial sector requires cybersecurity platforms to combine heterogeneous threat intelligence, domain knowledge, regulatory requirements, and operational context into a unified knowledge foundation. Within CyberAid, Task T2.3 addresses this need by defining a Security Knowledge Model (SKM) and a Cyber Threat Intelligence (CTI) Integration Framework capable of supporting semantic interoperability, knowledge contextualisation, and intelligent cybersecurity services across the project architecture. This deliverable presents the specification of the CyberAid SKM and the associated CTI Integration Framework developed during the M03–M09 period. The proposed approach establishes a common semantic foundation for representing cybersecurity concepts, financial-domain entities, regulatory obligations, and CTI artefacts, enabling consistent interpretation, correlation, and reuse of knowledge throughout the CyberAid ecosystem. Particular emphasis is placed on interoperability with widely adopted cybersecurity standards and structured threat intelligence representations, while maintaining sufficient flexibility to accommodate evolving financial-sector requirements and future platform extensions. The deliverable further defines the logical mechanisms for ingesting, normalising, enriching, and contextualising threat intelligence originating from multiple external sources. Rather than focusing on the implementation of specific technologies, it specifies the conceptual models, semantic relationships, metadata structures, governance principles, and interoperability mechanisms required to transform heterogeneous CTI into structured knowledge that can be consumed by downstream CyberAid components. In addition to the conceptual specification, D2.4 documents the modelling activities and technical investigations performed during the first implementation phase, including the exploration of ontology-based modelling, graph-oriented knowledge representation, and CTI management approaches used to validate the proposed specification. These activities provide an initial validation baseline while remaining independent of the final implementation choices that will be addressed in subsequent project deliverables.

Overall, D2.4 establishes the semantic and architectural foundations required to support cybersecurity knowledge management, CTI integration, explainable AI-driven analysis, and knowledge-based decision support within CyberAid, serving as a reference specification for the continued development, integration, and validation activities across the project.

TABLE OF CONTENTS

1.	Introduction	9
1.1	Scope of the Deliverable.....	9
1.2	Relation to Other Work Packages.....	9
1.3	Structure of the Deliverable.....	10
2.	Architectural Positioning and Alignment.....	10
2.1	Position within the CyberAid Reference Architecture and Requirements Baseline	10
2.2	Interaction with External Knowledge Sources.....	12
2.3	Possible Interfaces, Inputs and Outputs within CyberAid	13
3.	CyberAid Security Knowledge Model	14
3.1	SKM Definition, Purpose and Modelling Baseline	15
3.2	SKM Design Principles.....	16
3.3	Cybersecurity and CTI Semantic Domains	17
3.3.1	Semantic Organisation of the Cybersecurity Domain	18
3.3.2	Threat Representation Domain.....	20
3.3.3	Attack & Adversarial Behaviour Domain	21
3.3.4	Vulnerability & Exposure Domain.....	22
3.3.5	Detection, Observation & Incident Domain	24
3.3.6	Defensive Knowledge Domain.....	25
3.3.7	Cross-Domain Semantic Integration	27
3.3.8	Core SKM class catalogue.....	29
3.4	Financial, Risk and Regulatory Concepts.....	30
3.4.1	Financial Domain	31
3.4.2	Organisational Assets & Business Processes Domain	32
3.4.3	Risk Representation Domain	34
3.4.4	Regulatory & Compliance Domain.....	35
3.4.5	Cross-Domain Business Integration	37
3.5	Relationships and Semantic Associations	39
3.5.1	Semantic Relationship Model.....	39
3.5.2	Structural Relationships.....	41
3.5.3	Behavioural Relationships.....	42
3.5.4	Contextual Relationships	44
3.5.5	Temporal and Lifecycle Relationships.....	45
3.5.6	Cross-Domain Relationship Patterns	47
3.5.7	Relationship type catalogue	49
3.6	Metadata, Provenance and Governance Attributes	50

3.6.1 Metadata Model.....	51
3.6.2 Provenance and Trust.....	54
3.6.3 Knowledge Lifecycle.....	56
3.6.4 Governance and Access Control	58
3.6.5 Explainability and Auditability	61
4. CTI Integration Framework.....	62
4.1 Purpose and Scope of the CTI Integration Framework.....	64
4.2 CTI Source Categories and ISAC/FS-ISAC Considerations.....	65
4.3 CTI Ingestion, Normalization and Correlation	68
4.4 STIX-Compatible CTI-to-SKM Mapping Approach.....	71
4.5 Data Quality, Governance and Bidirectional Sharing.....	74
5. Knowledge Representation and Consumption	76
5.1 Knowledge Representation Models	76
5.2 Semantic and LLM-Ready Representation	78
5.3 Knowledge Consumption Across CyberAid	80
6. Conclusions And Next Steps	83
References	85

LIST OF FIGURES

Figure 1 - Position of SKM and CTI Integration Framework.....	12
Figure 2 - Recommended Logical Processing Flow	71

LIST OF TABLES

Table 1 - Design Principles.....	16
Table 2 - SKM Semantic Domain Overview	18
Table 3 - Class Catalogue.....	29
Table 4 - Relationship Catalogue	49
Table 5 - Metadata and Provenance.....	51
Table 6 - Confidence Scale	55
Table 7 - Classification.....	59
Table 8 - TLP Mapping.....	60
Table 9 - Representative CTI Source Categories.....	67
Table 10 - Representative STIX-to-SKM Mapping	73
Table 11 - Domain to consumer map.....	81

1. INTRODUCTION

1.1 Scope of the Deliverable

This deliverable specifies the Security Knowledge Model (SKM) and the Cyber Threat Intelligence (CTI) Integration Framework developed within Task T2.3 of the CyberAid project. Its primary objective is to establish a common semantic foundation for representing, organising, and contextualising cybersecurity knowledge relevant to the financial sector, enabling consistent interpretation and interoperability across the CyberAid ecosystem. By defining a structured knowledge model and the corresponding CTI integration mechanisms, D2.4 contributes to the harmonisation of heterogeneous cybersecurity information originating from multiple internal and external sources, while supporting subsequent architectural, analytical, and operational activities throughout the project. The proposed specification extends the requirements baseline established in D2.1 by translating the identified functional and technical requirements into a semantic knowledge representation capable of supporting financial-sector cybersecurity scenarios. In particular, the deliverable defines the concepts, relationships, metadata structures, and governance principles required to represent cybersecurity entities, financial-domain assets, regulatory obligations, and threat intelligence in a coherent and interoperable manner. Furthermore, it specifies the logical mechanisms required to ingest, normalise, correlate, enrich, and contextualise CTI from heterogeneous intelligence sources before integrating it into the SKM. Rather than prescribing a specific technological implementation, D2.4 defines the conceptual models and interoperability principles that will guide the subsequent implementation activities within the CyberAid platform. Candidate technologies explored during the first project period, including ontology modelling environments, CTI management platforms, graph-based representations, and semantic knowledge processing approaches, are considered solely as indicative validation environments used to assess the feasibility of the proposed specification. Consequently, implementation-specific decisions, technology selection, and platform integration activities remain outside the scope of this deliverable and will be addressed in subsequent technical deliverables.

This deliverable consolidates the initial semantic modelling and CTI integration specification developed during the M03–M09 period. It establishes a conceptual baseline for subsequent implementation, technical validation, integration and refinement activities. Evidence from prototype implementation and operational validation will be documented in the corresponding subsequent technical deliverables.

1.2 Relation to Other Work Packages

Task T2.3 constitutes one of the core knowledge engineering activities within CyberAid, providing the semantic and cyber threat intelligence foundations upon which several technical components of the project are built. While D2.1 establishes the cybersecurity requirements and domain-specific needs of the financial sector, D2.4 translates these requirements into a structured semantic specification by defining the SKM and the CTI Integration Framework. Consequently, this deliverable acts as an intermediate specification layer between the project requirements and the subsequent architectural, implementation, and validation activities. Within WP2, the proposed semantic models complement the CyberAid Reference Architecture by specifying how cybersecurity concepts, financial-domain entities, regulatory knowledge, and cyber threat intelligence can be represented, organised, and exchanged in a consistent manner. Rather than introducing new architectural components, D2.4 provides the semantic structures required to support interoperability among the services and data flows defined within the overall platform architecture.

D2.2 defines the CyberAid reference architecture and positions the SKM and CTI capabilities within the Knowledge and Data Fabric. D2.4 refines this architectural area by specifying its semantic domains, ontologies and taxonomies, relationship model, metadata, provenance, governance attributes and CTI integration mechanisms. D2.5 consumes these semantic specifications through the LLM orchestration and agentic-analysis design, including governed knowledge retrieval and evidence-grounded analytical outputs. The division of responsibility is explicit: D2.4 specifies the semantic knowledge and CTI integration contracts; D2.5 specifies the LLM agent architecture, workflows, orchestration logic, Security Context

processing and runtime knowledge access; the corresponding WP4 tasks specify operational monitoring, event generation, alerting and notification mechanisms.

Beyond WP2, the deliverable establishes a common knowledge foundation that can be exploited by multiple technical work packages. The SKM provides structured knowledge representations that can support cybersecurity modelling, knowledge-driven analytics, graph-based reasoning, and AI-assisted decision support, while the CTI Integration Framework specifies the mechanisms required to transform heterogeneous threat intelligence into interoperable and reusable knowledge. These capabilities are expected to facilitate the development of advanced cybersecurity functionalities across CyberAid while ensuring semantic consistency between independently developed components. Furthermore, the specification presented in this deliverable supports the progressive evolution of CyberAid throughout the project lifecycle. As additional cybersecurity services, knowledge sources, pilot requirements, and validation scenarios become available, the SKM and the CTI Integration Framework are expected to evolve accordingly while preserving a common semantic backbone across the platform. In this context, D2.4 serves not only as a technical specification for the current reporting period but also as a reference document that guides future architecture refinement, component integration, and pilot validation activities.

1.3 Structure of the Deliverable

This deliverable is organised into six main sections that progressively introduce the semantic foundations of the CyberAid SKM and the associated CTI Integration Framework. Following the introductory chapter, the document positions the proposed approach within the CyberAid Reference Architecture and explains its relationship with the project requirements, external knowledge sources, and the wider project ecosystem. The core chapters subsequently define the SKM, the CTI Integration Framework, and the different knowledge representation approaches that support interoperability, semantic consistency, and the utilisation of cybersecurity knowledge across CyberAid components. The final section summarises the main contributions of the deliverable and outlines the next steps required to further refine the specification and support its evolution throughout the implementation and integration phases of the project. References to relevant standards, regulatory frameworks, scientific literature, and project deliverables are provided throughout the document where appropriate to support the proposed concepts and design decisions.

2. ARCHITECTURAL POSITIONING AND ALIGNMENT

The SKM and the CTI Integration Framework form the semantic foundation that enables the CyberAid platform to integrate, contextualise, and exchange cybersecurity knowledge across its architectural components. Their effectiveness, however, depends not only on the definition of semantic concepts and knowledge structures, but also on their alignment with the overall platform architecture, the external knowledge ecosystem, and the operational requirements of the participating services. This chapter positions the proposed semantic specification within the CyberAid architecture by describing its architectural role, its interaction with external knowledge sources, and its logical interfaces with the remaining platform capabilities. Establishing this architectural context provides the basis for the detailed specification of the SKM and the CTI Integration Framework presented in the subsequent chapters.

2.1 Position within the CyberAid Reference Architecture and Requirements Baseline

The CyberAid platform has been designed as an intelligent cybersecurity ecosystem capable of supporting financial-sector organisations throughout the complete cybersecurity lifecycle, from cyber threat intelligence collection and knowledge management to intelligent analysis, risk assessment, regulatory compliance, and operational decision support. Achieving these objectives requires the coordinated interaction of multiple architectural components that operate on heterogeneous information originating from diverse technical, operational, regulatory, and organisational domains. As these information domains differ significantly in structure, semantics, granularity, reliability, and update frequency, effective interoperability cannot rely solely on syntactic data exchange mechanisms. Instead, CyberAid requires a

common semantic foundation capable of ensuring that cybersecurity knowledge is represented, interpreted, exchanged, and reused consistently across all platform capabilities. The SKM fulfils this role by establishing the semantic backbone of the CyberAId ecosystem. The SKM is a technology-independent semantic specification rather than an isolated software module or data repository. It defines the concepts, relationships and governance attributes that implementations shall preserve, while specific storage and processing technologies remain implementation choices.

The architectural role of the SKM becomes particularly important when considering the diversity of information processed by CyberAId. Cyber threat intelligence, vulnerability disclosures, attack knowledge, regulatory obligations, operational security events, financial-sector assets, business processes, and risk-related information are generated independently and follow different representation models. Without a common semantic layer, each platform component would be required to interpret and correlate these information sources independently, leading to duplicated knowledge, inconsistent interpretations, and reduced interoperability. The SKM addresses this challenge by providing a shared conceptual representation through which heterogeneous knowledge can be organised according to common semantics while preserving its original context, provenance, confidence, temporal validity, and domain-specific characteristics. Within the CyberAId Reference Architecture defined in D2.2, the SKM and CTI integration capabilities are positioned in the Knowledge and Data Fabric, where they provide the common semantic structures required to organise, contextualise and retrieve cybersecurity knowledge independently of its origin. Information originating from external intelligence providers, internal platform services, regulatory repositories, financial-domain sources, and operational monitoring components is progressively transformed into structured knowledge that can be shared across the platform. The semantic structures defined by the SKM allow individual knowledge elements to be connected through explicit relationships, enriched with contextual metadata, and maintained according to common governance principles. Consequently, the knowledge infrastructure evolves from a collection of heterogeneous datasets into an interconnected semantic knowledge ecosystem capable of supporting higher-level analytical processes. This semantic foundation directly supports the operation of the platform's intelligent capabilities. AI-assisted services, domain-specific agents, cybersecurity monitoring components, decision support mechanisms, and reporting functionalities require access not only to isolated data but also to the relationships that connect cybersecurity entities, financial assets, regulatory obligations, historical incidents, vulnerabilities, attack techniques, and defensive controls. By organising this information within a common semantic framework, the SKM enables these components to exploit contextual knowledge rather than individual datasets. This approach facilitates explainable reasoning, improves the consistency of analytical results, supports graph-oriented knowledge exploration, and promotes knowledge reuse across different cybersecurity services without introducing dependencies between their internal implementations. The architectural positioning of the SKM also establishes a clear separation of responsibilities between knowledge representation and knowledge processing. The semantic model specifies the concepts that exist within the CyberAId knowledge ecosystem, the relationships that connect them, and the metadata required to ensure traceability, provenance, and governance. Conversely, the mechanisms responsible for acquiring cyber threat intelligence, integrating external information sources, executing reasoning processes, performing AI inference, or presenting cybersecurity insights remain independent architectural concerns. This separation promotes modularity, facilitates future evolution of the platform, and enables individual technological components to be replaced or extended without affecting the underlying semantic specification.

Another important architectural characteristic concerns extensibility. Cybersecurity knowledge is inherently dynamic, with new attack techniques, threat actors, regulatory obligations, financial technologies, and operational requirements continuously emerging. Consequently, the SKM has been designed as an extensible semantic specification capable of accommodating additional concepts, relationships, and knowledge domains without requiring structural redesign. This capability is particularly important within CyberAId, where new pilot requirements, additional CTI providers, AI capabilities, and regulatory developments are expected to become available throughout the project. Maintaining a stable semantic core while allowing controlled knowledge evolution ensures long-term interoperability and

facilitates the incremental expansion of the platform. Finally, the architectural role of the SKM extends beyond the boundaries of individual CyberAid services by establishing a common semantic language that supports collaboration between independently developed components. The proposed specification enables cybersecurity knowledge to be exchanged according to shared conceptual definitions rather than technology-specific data structures, reducing semantic ambiguity while facilitating integration across the platform. As a result, the SKM represents one of the enabling foundations of the CyberAid architecture, providing the semantic consistency upon which CTI integration, knowledge-driven analytics, intelligent orchestration, explainable AI, and cybersecurity decision support can be progressively realised throughout the subsequent implementation and validation phases of the project.




Figure 1 - Position of SKM and CTI Integration Framework

As shown in Figure 1, the SKM and CTI Integration Framework are positioned within the Knowledge and Data Fabric defined by the D2.2 Reference Architecture. External CTI, vulnerability, regulatory and organisational sources provide controlled inputs, while authorised CyberAid capabilities consume governed semantic knowledge and evidence.

2.2 Interaction with External Knowledge Sources

The SKM is intended to operate within a dynamic cybersecurity ecosystem where knowledge continuously evolves as new threats, vulnerabilities, technologies, regulations, and operational practices emerge. Consequently, the semantic model is designed to interact with a broad spectrum of external knowledge sources that collectively provide the information required to maintain an accurate, contextualised, and continuously evolving representation of the cybersecurity landscape. Rather than relying on individual repositories or predefined datasets, the proposed specification accommodates multiple complementary knowledge domains, each contributing distinct semantic information to the overall knowledge ecosystem.

CTI represents one of the principal knowledge domains considered by the SKM. CTI contributes structured and contextual information describing threat actors, adversary campaigns, attack techniques, malware

families, indicators of compromise, observed incidents, tactics, procedures, and emerging threat trends. The value of this information extends beyond the description of isolated cyber events, enabling the semantic model to capture relationships between adversaries, attack behaviours, targeted assets, and defensive measures. As threat intelligence continuously evolves, its incorporation into the knowledge ecosystem ensures that the semantic representation remains relevant to the changing cybersecurity landscape.

The semantic model is further enriched through interaction with vulnerability and attack knowledge repositories. These knowledge domains provide structured information regarding software vulnerabilities, configuration weaknesses, affected products, exploitation conditions, attack patterns, mitigation strategies, and security controls. Such information enables the representation of explicit relationships between vulnerabilities, assets, attack techniques, defensive mechanisms, and associated risks, facilitating a more comprehensive understanding of cybersecurity exposures and their potential operational impact. The combination of vulnerability intelligence and attack knowledge therefore complements traditional CTI by providing the technical context required for knowledge correlation and security analysis.

Beyond purely technical cybersecurity information, the SKM incorporates external knowledge describing regulatory, organisational, and financial-sector contexts. Regulatory frameworks, technical standards, supervisory guidance, and compliance obligations provide the governance perspective required to associate cybersecurity activities with legal, operational, and organisational requirements. In parallel, financial-domain knowledge contributes information describing business services, financial assets, payment infrastructures, digital banking ecosystems, operational processes, and sector-specific terminology. Integrating these complementary knowledge domains enables cybersecurity events to be interpreted not only according to their technical characteristics but also according to their potential operational, business, and regulatory implications within the financial sector.

Recognising that cybersecurity knowledge originates from diverse communities and information-sharing ecosystems, the proposed specification has been designed to remain independent of individual knowledge providers while supporting the incorporation of information originating from trusted public, private, and collaborative environments. These may include cyber threat intelligence repositories, vulnerability databases, attack knowledge bases, regulatory publications, Information Sharing and Analysis Centers (ISACs), Financial Services Information Sharing and Analysis Centers (FS-ISACs), and other trusted intelligence-sharing initiatives. Their contribution depends on factors such as accessibility, licensing conditions, information handling restrictions, organisational participation, and data-sharing agreements. Accordingly, the SKM adopts a provider-agnostic approach, ensuring that the semantic representation remains stable even as individual knowledge sources evolve, are replaced, or become unavailable.

The interaction with external knowledge sources establishes the knowledge ecosystem upon which the SKM is built. While these sources provide the raw knowledge required to describe the cybersecurity domain, they differ significantly in terms of structure, semantics, quality, granularity, and update frequency. The proposed specification therefore treats external knowledge as complementary and heterogeneous inputs that collectively enrich the semantic model while preserving their provenance, contextual information, and domain-specific characteristics. The mechanisms responsible for acquiring, transforming, validating, and semantically integrating this information are specified separately within the CTI Integration Framework, thereby maintaining a clear separation between the definition of knowledge sources and the processes through which such knowledge becomes part of the CyberAid semantic ecosystem.

2.3 Possible Interfaces, Inputs and Outputs within CyberAid

The SKM establishes the common semantic interaction layer through which cybersecurity knowledge is exchanged, interpreted, and reused across the CyberAid ecosystem. Rather than acting as a conventional data repository, the model defines the semantic interface that enables heterogeneous platform capabilities to communicate through a common understanding of cybersecurity concepts, relationships, contextual information, and governance attributes. This interaction model ensures that knowledge remains consistent and interoperable regardless of the technologies, programming languages, storage mechanisms, or analytical services employed by individual platform components. The interaction model is based on the

principle that knowledge producers and knowledge consumers remain decoupled from one another through a common semantic representation. Components responsible for monitoring, intelligence acquisition, vulnerability analysis, compliance assessment, or operational data collection are not expected to exchange proprietary data structures directly. Instead, the information they generate is represented according to the concepts and relationships defined by the SKM, allowing knowledge originating from different domains to be interpreted consistently and reused throughout the platform. This semantic decoupling promotes interoperability, reduces integration complexity, and supports the independent evolution of platform components without affecting the overall knowledge ecosystem.

Any realisation of the SKM shall support the logical partitioning of retrievable knowledge into identifiable knowledge scopes, the association of entities and relationships with one or more named semantic domains or graph scopes, and the assignment of classification and handling attributes to every governed knowledge element. These properties constitute semantic metadata through which authorised consumers can request an appropriately delimited knowledge view. The binding of these properties to a Security Context, the technical retrieval interface and the enforcement of access decisions are defined by the corresponding architectural and orchestration components.

The SKM accommodates knowledge originating from multiple complementary domains, each contributing different perspectives to the overall cybersecurity representation. Cyber threat intelligence, vulnerability information, attack knowledge, regulatory and compliance information, financial-sector context, operational security events, and domain-specific business information collectively constitute the semantic inputs of the model. Rather than treating these domains as isolated datasets, the proposed specification enables them to be represented within a unified semantic context where relationships between technical, organisational, operational, and regulatory knowledge can be explicitly established. This integrated representation allows individual knowledge domains to reinforce one another while preserving their provenance, confidence, temporal validity, and domain-specific characteristics. Once represented according to the common semantic specification, cybersecurity knowledge becomes available to the different architectural capabilities of the CyberAid platform. The semantic representation provides contextual information that can support intelligent cybersecurity monitoring, knowledge-driven analytics, graph exploration, AI-assisted reasoning, regulatory compliance assessment, risk analysis, incident investigation, and decision support. By exposing knowledge through common semantic structures rather than technology-specific interfaces, the SKM enables these capabilities to exploit contextual relationships instead of isolated information, improving the consistency and explainability of analytical outcomes.

The proposed interaction model has also been designed to support the continuous evolution of the CyberAid platform. As additional intelligence sources, financial-domain knowledge, regulatory requirements, pilot-specific assets, AI services, or analytical components become available, they can be incorporated into the semantic ecosystem through the common conceptual structures defined by the model rather than through individual point-to-point integrations. This approach promotes scalability, facilitates knowledge reuse, and preserves semantic consistency as the platform progressively expands throughout the project lifecycle. Finally, the SKM establishes the architectural boundary between semantic representation and knowledge processing. The model specifies the semantic structures required to organise and exchange cybersecurity knowledge, while the mechanisms responsible for knowledge acquisition, transformation, correlation, reasoning, storage, and presentation remain independent architectural concerns. This separation ensures that the semantic specification remains stable despite future technological evolution and provides a robust foundation for the detailed SKM and CTI Integration Framework specified in the following chapters. Runtime persistence, retrieval and orchestration technology choices are specified separately in D2.5.

3. CYBERAID SECURITY KNOWLEDGE MODEL

The SKM constitutes the semantic core of the CyberAid knowledge ecosystem and provides the common conceptual foundation through which cybersecurity knowledge is represented, contextualised, and exchanged across the platform. Its purpose is not limited to organising information originating from

different cybersecurity domains, but rather to establish a unified semantic abstraction capable of connecting technical, operational, organisational, financial, and regulatory knowledge within a coherent representation. By defining common semantic concepts, relationships, contextual information, and governance attributes, the SKM enables heterogeneous knowledge to be interpreted consistently regardless of its origin, structure, or intended consumer. The proposed model has been developed according to knowledge engineering principles that prioritise semantic interoperability, extensibility, traceability, and knowledge reuse over technology-specific implementation choices. Consequently, the SKM should be regarded as a conceptual specification rather than a software artefact or a predefined ontology implementation. Its conceptual structures are intentionally independent of ontology languages, graph technologies, databases, or knowledge management platforms, allowing different implementation approaches to realise the same semantic model while preserving a common interpretation of cybersecurity knowledge throughout the CyberAid ecosystem.

The following sections specify the SKM's modelling baseline, design principles, semantic domains, relationships and governance attributes.

3.1 SKM Definition, Purpose and Modelling Baseline

The SKM defines the common concepts, relationships and governance attributes used to represent cybersecurity, financial, organisational and regulatory knowledge across CyberAid. It harmonises heterogeneous source models at the semantic level while remaining independent of storage, ontology and processing technologies.

The proposed specification adopts a knowledge-centric modelling philosophy in which cybersecurity knowledge is represented as an interconnected semantic ecosystem rather than as a collection of independent datasets. Every knowledge entity is expected to contribute to a broader semantic context by maintaining explicit relationships with other entities that collectively describe the operational cybersecurity environment. This modelling philosophy reflects the observation that cybersecurity events rarely exist in isolation. Vulnerabilities affect assets, assets support business services, business services are subject to regulatory obligations, attack techniques exploit vulnerabilities, threat actors execute attack campaigns, security controls mitigate identified risks, and operational events provide evidence of ongoing activities. Representing these dependencies explicitly enables the semantic model to capture the operational context in which cybersecurity knowledge is created and utilised.

Every SKM entity, relationship, source artefact and provenance record shall be independently addressable through a globally unique and resolvable identifier. The canonical identifier represents the semantic identity of the knowledge element and shall remain stable across re-ingestion, enrichment, validation, deduplication, lifecycle transitions and changes to its physical storage location. Updates shall create a new version of the identified element rather than a new semantic identity, unless the information represents a genuinely distinct object or assertion.

To achieve this objective, the SKM adopts a domain-oriented conceptual organisation. Instead of describing the cybersecurity domain through a single hierarchical taxonomy, the model is structured around multiple complementary semantic domains, each representing a distinct perspective of the operational environment. These domains include cybersecurity knowledge, cyber threat intelligence, financial-sector knowledge, organisational and operational context, governance and compliance information, and metadata supporting provenance and lifecycle management. Although each domain maintains its own conceptual boundaries, they are intentionally interconnected through explicit semantic relationships that allow knowledge to be progressively contextualised and correlated across technical and business perspectives. The detailed specification of these semantic domains is presented in the following sections of this chapter. A fundamental characteristic of the CyberAid SKM is its technology-independent nature. The specification intentionally avoids prescribing implementation technologies, storage mechanisms, ontology languages, graph databases, or knowledge management platforms. Instead, it defines the conceptual semantics that every implementation is expected to preserve regardless of the selected technological solution. This separation between semantic specification and implementation allows the model to be realised through

ontology-based representations, knowledge graphs, semantic databases, document-oriented repositories, or hybrid knowledge management solutions without affecting the consistency of the underlying conceptual model. It also enables future technological evolution while preserving semantic interoperability across successive platform versions. The modelling baseline further incorporates established knowledge engineering principles to maximise interoperability with existing cybersecurity ecosystems and recognised domain standards. The specification promotes semantic consistency through explicit concept definitions, relationship semantics, controlled contextualisation, metadata governance, and clear separation between conceptual entities and implementation artefacts. The SKM reuses the STIX 2.1 object, relationship, data-marking and confidence semantics for CTI exchanged in STIX form [4]. It maps its threat and adversarial-behaviour concepts to MITRE ATT&CK [6], its vulnerability concepts to CVE [7] and NVD [8] records and CVSS assessments [9], and its weakness and attack-pattern concepts to CWE [10] and CAPEC [11]. The authoritative identifiers and version information of these external resources shall be preserved. Rather than reproducing these resources, the SKM provides the semantic framework through which they can be interpreted and interconnected within a common operational context. Another important modelling objective concerns extensibility. Cybersecurity continuously evolves through the emergence of new threat actors, attack techniques, defensive mechanisms, regulatory obligations, digital financial services, and operational practices. Consequently, the SKM is specified as an extensible semantic framework capable of incorporating new knowledge domains, concepts, relationships, and contextual information without requiring structural redesign. This extensibility enables the model to accommodate future intelligence sources, additional pilot requirements, evolving regulatory frameworks, and new CyberAid capabilities while preserving the integrity of the existing semantic structures.

Finally, the SKM establishes the conceptual baseline upon which the remaining specification is progressively developed. The following sections refine this baseline by defining the design principles governing the model, the semantic concepts representing the cybersecurity and financial domains, the relationships that connect these concepts, and the metadata required to support provenance, governance, traceability, and long-term knowledge evolution. Together, these elements constitute the semantic specification that underpins knowledge representation, CTI integration, and knowledge-driven cybersecurity services within the CyberAid ecosystem.

3.2 SKM Design Principles

The design of the CyberAid SKM is governed by a set of knowledge engineering principles intended to ensure that the semantic model remains consistent, interoperable, extensible, and capable of supporting the diverse cybersecurity services envisioned within the CyberAid ecosystem. Rather than representing isolated design objectives, these principles collectively define the modelling philosophy upon which the semantic specification is constructed and guide the organisation of concepts, relationships, metadata, and governance attributes throughout the model.

Table 1 - Design Principles

Principle	Specification statement	Required consequence
Semantic consistency	Each concept and relationship has one defined meaning within its semantic domain.	Implementations shall validate element types, attributes and relationship semantics against the SKM catalogues.
Semantic interoperability	External information is mapped to common SKM concepts while retaining its source representation.	External identifiers, mappings and mapping provenance shall be preserved.

Principle	Specification statement	Required consequence
Contextualisation	Knowledge is associated with its operational, organisational, temporal, financial and regulatory context.	Retrievable elements shall expose sufficient context for their correct interpretation.
Modularity and domain separation	Concepts are organised into bounded semantic domains connected through explicit cross-domain relationships.	Extensions shall identify their owning domain and shall not silently redefine concepts belonging to another domain.
Technology independence	Consumers depend on semantic contracts rather than physical persistence technologies.	Storage changes shall not modify canonical identifiers, concept semantics or required metadata.
Controlled extensibility	New concepts and relationships may be added without silently changing existing semantics.	Extensions shall be versioned and reviewed for semantic and backward compatibility.
Provenance and governance by design	Origin, transformation, confidence, classification and lifecycle information accompany governed knowledge.	Required provenance and governance fields shall be validated before knowledge is made available.
Reusability	The same knowledge element may support multiple authorised analytical purposes.	Reuse shall preserve identity, provenance, classification and source-imposed constraints.

These principles constitute requirements on any implementation claiming conformance with the SKM specification. They do not prescribe the database, ontology language, retrieval service or orchestration technology through which conformance is achieved.

3.3 Cybersecurity and CTI Semantic Domains

The SKM organises cybersecurity knowledge into bounded but interconnected semantic domains. Each domain defines a coherent set of concepts, while typed cross-domain relationships connect technical, operational, organisational, financial, risk and regulatory knowledge. The following subsections specify the purpose, principal entities and integration role of each domain. The semantic domains specified in this section should not be interpreted as isolated repositories of information. Instead, they constitute complementary views of a common knowledge space in which every semantic entity contributes to the description of the same operational environment. Information represented within one domain is expected to establish explicit semantic relationships with entities belonging to other domains, enabling cybersecurity knowledge to evolve from isolated technical observations towards comprehensive representations of threats, organisational impact, regulatory relevance, and defensive response. This interconnected organisation forms the basis for semantic interoperability, graph-based reasoning, contextual analytics, and AI-assisted knowledge exploitation throughout the CyberAid platform.

The following subsections progressively define the principal cybersecurity and CTI semantic domains that constitute the conceptual backbone of the SKM. Each domain is specified according to its purpose, semantic scope, principal entities, contribution to the overall knowledge ecosystem, and relationships with the remaining domains. Together, these domains establish a comprehensive semantic architecture capable of

supporting the representation, contextualisation, and continuous evolution of cybersecurity knowledge across the CyberAId ecosystem.

Table 2 - SKM Semantic Domain Overview

Semantic domain	Primary scope	Representative concepts
Threat Representation	Threat intelligence and adversary context	ThreatActor, Campaign, Indicator, CTI Report
Attack and Adversarial Behaviour	Adversary tactics, techniques and attack execution	AttackTechnique, Malware, Tool, AttackPath
Vulnerability and Exposure	Technical weaknesses and affected resources	Vulnerability, Weakness, Product, Asset
Detection, Observation and Incident	Operational evidence and analytical results	Observation, Detection, Finding, Incident
Defensive Knowledge	Preventive, detective and responsive measures	SecurityControl, Mitigation, ResponseAction
Financial	Financial-sector services and infrastructures	FinancialService, PaymentInfrastructure, FinancialProduct
Organisational Assets and Business Processes	Organisational context and dependencies	Organisation, Asset, BusinessProcess, Service
Risk Representation	Cyber-risk conditions and assessments	Risk, Likelihood, Impact, RiskLevel
Regulatory and Compliance	Obligations, controls and supporting evidence	RegulatoryFramework, RegulatoryRequirement, ComplianceEvidence

These semantic domains are complemented by the cross-cutting metadata, provenance, lifecycle and governance attributes specified in Section 3.6.

3.3.1 Semantic Organisation of the Cybersecurity Domain

The CyberAId SKM organises cybersecurity knowledge through a structured semantic architecture designed to represent the operational cybersecurity ecosystem as a collection of interconnected knowledge domains rather than independent information repositories. This modelling approach reflects the multidisciplinary nature of cybersecurity, where threats, vulnerabilities, attacks, assets, regulations, operational processes, and defensive activities continuously influence one another throughout the cybersecurity lifecycle. Consequently, the proposed semantic organisation aims to represent not only individual cybersecurity entities but also the contextual relationships that determine their operational significance.

Existing resources generally cover bounded cybersecurity domains: CWE catalogues software and hardware weaknesses, CAPEC catalogues attack patterns, MITRE ATT&CK describes adversary behaviour, and NVD publishes vulnerability records [6], [8], [10], [11]. The SKM complements these resources through a domain-oriented model that links their concepts with operational, organisational, financial, risk and regulatory knowledge. Each semantic domain represents a distinct aspect of cybersecurity knowledge while remaining explicitly interconnected with the remaining domains through well-defined semantic relationships. This separation enables the model to preserve conceptual clarity, simplify knowledge evolution, and facilitate the progressive extension of the semantic ecosystem without introducing unnecessary dependencies between unrelated concepts. The semantic organisation further distinguishes between foundational knowledge and contextual knowledge. Foundational knowledge represents relatively stable concepts that describe the cybersecurity domain itself, including threats, vulnerabilities, attack techniques, security controls, assets, and organisational entities. Contextual knowledge enriches these concepts by introducing operational observations, intelligence reports, incidents, regulatory obligations, business processes, historical evidence, and additional metadata that describe how foundational entities interact within specific operational environments. This distinction enables the SKM to maintain a stable conceptual backbone while supporting the continuous incorporation of dynamic cybersecurity information.

From a modelling perspective, every semantic domain contributes complementary knowledge that becomes progressively interconnected through explicit semantic relationships. Rather than treating cyber threat intelligence, vulnerability information, financial-domain context, or operational observations as independent datasets, the SKM considers them complementary semantic perspectives describing different dimensions of the same cybersecurity environment. This multidimensional representation enables analytical processes to move seamlessly between technical observations, operational context, business relevance, and regulatory implications while maintaining semantic consistency across the entire knowledge ecosystem. The proposed semantic organisation also introduces different levels of abstraction. Low-level concepts represent concrete cybersecurity observations, such as vulnerabilities, malware samples, indicators of compromise, software assets, or security events. Higher abstraction levels capture operational concepts including attack campaigns, business services, organisational units, risk scenarios, regulatory obligations, defensive capabilities, and strategic cybersecurity objectives. Establishing explicit relationships between these abstraction levels enables the model to support both detailed technical analysis and higher-level organisational reasoning without duplicating knowledge or introducing multiple semantic representations of the same operational reality.

Another distinguishing characteristic concerns lifecycle representation. Cybersecurity knowledge is not considered static but continuously evolves throughout the lifecycle of cyber threats and defensive activities. Information initially introduced as threat intelligence may later become operational evidence, contribute to incident investigations, trigger defensive actions, influence risk assessments, and ultimately enrich organisational knowledge. The semantic organisation therefore supports the progressive enrichment of knowledge throughout its operational lifecycle, preserving historical context while enabling new relationships and contextual information to be incorporated over time. This semantic architecture has also been designed to facilitate knowledge reuse across independent CyberAid services. Since all semantic domains share common modelling principles and governance rules, cybersecurity knowledge represented within one domain can be reused by multiple analytical capabilities without requiring semantic transformation. Threat intelligence may therefore support attack analysis, vulnerability prioritisation, incident investigation, compliance assessment, AI-assisted reasoning, graph analytics, or decision support while maintaining a single conceptual representation throughout the platform.

Finally, the semantic organisation establishes the conceptual framework upon which the remaining cybersecurity domains are specified. Each domain captures a specific operational perspective of cybersecurity knowledge, while their integration provides the complete semantic representation required by the CyberAid ecosystem. The following subsections progressively define these domains, beginning with the representation of threats and adversarial entities before introducing attack behaviour, vulnerabilities, operational observations, defensive knowledge, and finally the semantic mechanisms that interconnect the entire knowledge ecosystem.

3.3.2 Threat Representation Domain

The Threat Representation Domain constitutes the strategic layer of the CyberAid SKM, providing the semantic structures required to describe the entities responsible for malicious cyber activities and the broader threat landscape in which these activities evolve. Unlike lower-level operational observations, which describe individual technical events or security artefacts, the Threat Representation Domain focuses on modelling the origin, intent, capabilities, objectives, and evolution of adversarial behaviour. Its primary objective is to establish a coherent semantic representation of the threat ecosystem, enabling CyberAid to interpret cybersecurity observations within the broader context of adversarial operations rather than as isolated technical occurrences. Within the proposed semantic architecture, threats are not represented as independent entities but as components of an interconnected operational ecosystem. A threat may emerge from one or more threat actors operating individually or collectively through organised groups, intrusion sets, campaigns, or financially motivated criminal organisations. These adversarial entities are characterised not only by their technical capabilities but also by their strategic objectives, motivations, targeting preferences, operational maturity, geographical relevance, historical activities, and observed behavioural patterns. Representing these characteristics explicitly allows the SKM to capture the strategic dimension of cybersecurity and to relate individual technical observations to broader adversarial activities.

The Threat Representation Domain therefore introduces a semantic distinction between threat entities and threat activities. Threat entities describe the actors, organisations, or groups responsible for malicious operations, while threat activities represent the campaigns, objectives, preparation stages, and strategic initiatives through which these entities pursue their goals. This distinction allows the model to represent the long-term evolution of adversarial operations without associating every technical observation directly with an individual threat actor. Consequently, campaigns may evolve over time, involve multiple attack techniques, target different financial organisations, or reuse common operational infrastructures while remaining semantically connected to the responsible adversarial entities.

Another important modelling objective concerns the representation of adversarial intent. Cybersecurity events often provide evidence of technical activity without immediately revealing the strategic objectives behind the attack. The proposed semantic model therefore separates observable malicious behaviour from inferred adversarial objectives, allowing analytical processes to distinguish between factual observations and higher-level threat assessments. Concepts such as financial fraud, credential theft, ransomware deployment, service disruption, espionage, supply-chain compromise, or data exfiltration are consequently represented as semantic characteristics associated with adversarial activities rather than with isolated technical events. This distinction enables the progressive refinement of threat knowledge as additional intelligence becomes available. The Threat Representation Domain also captures the operational characteristics of threat actors and campaigns through contextual semantic attributes. These attributes may include motivation, capability level, sophistication, confidence, targeting sector, geographical focus, operational status, attribution confidence, associated infrastructure, historical observations, and temporal validity. Such contextual information enables CyberAid to preserve not only the identity of adversarial entities but also the circumstances under which they have been observed, thereby supporting longitudinal threat analysis and the continuous evolution of threat intelligence.

Since CyberAid targets the financial sector, particular emphasis is placed on representing threats according to their potential impact on financial services and digital banking infrastructures. Threat actors targeting payment systems, online banking platforms, financial APIs, customer authentication services, trading platforms, digital wallets, interbank communication infrastructures, or critical financial business processes may therefore be semantically associated with the corresponding organisational assets and business services defined in the financial domain of the SKM. This contextual representation enables strategic threat information to be interpreted according to the operational realities of financial institutions rather than solely through generic cybersecurity terminology. The Threat Representation Domain has also been designed to support continuous intelligence evolution. Threat intelligence rarely represents definitive knowledge; instead, it evolves as additional indicators, investigations, intelligence reports, or collaborative information-sharing activities become available. Consequently, threat entities are expected to evolve throughout their lifecycle by accumulating new contextual information, behavioural observations,

attribution evidence, and relationships with campaigns, attack techniques, vulnerabilities, or defensive actions. The semantic model therefore supports incremental knowledge enrichment while preserving historical context and provenance, allowing threat intelligence to mature progressively rather than being repeatedly replaced. From an architectural perspective, the Threat Representation Domain establishes the semantic entry point for subsequent cybersecurity domains. Threat entities provide the strategic context from which attack behaviour originates, vulnerabilities become exploited, operational events emerge, and defensive mechanisms are activated. Consequently, this domain forms the highest abstraction level of the cybersecurity semantic architecture and provides the contextual knowledge required to interpret all downstream cybersecurity observations represented within the SKM.

3.3.3 Attack & Adversarial Behaviour Domain

The Attack & Adversarial Behaviour Domain defines the semantic structures through which malicious activities, offensive techniques, attack progression, and adversarial behaviour are represented within the CyberAid SKM. Building upon the strategic context established by the Threat Representation Domain, this semantic domain captures how adversaries operate, which methods they employ, how attacks evolve, and how malicious activities materialise throughout the different phases of an attack lifecycle. Its purpose is to provide a structured representation of adversarial behaviour that enables CyberAid to interpret technical observations as components of coordinated attack operations rather than isolated cybersecurity events. Unlike the Threat Representation Domain, which focuses on the strategic characteristics of adversaries, the Attack & Adversarial Behaviour Domain represents the operational execution of malicious activities. Every attack is considered the result of one or more adversarial decisions that materialise through a sequence of techniques, procedures, tools, exploits, and malicious actions. The semantic model therefore distinguishes between who performs an attack and how the attack is executed, allowing the behaviour of different adversaries to be analysed independently from their identity while simultaneously preserving the semantic links between strategic intent and operational execution.

The proposed specification models attacks as dynamic processes rather than static events. Modern cyber-attacks evolve through multiple operational phases that may include reconnaissance, initial access, execution, persistence, privilege escalation, defence evasion, credential access, lateral movement, discovery, command-and-control communications, collection, exfiltration, impact, and post-compromise activities. The semantic model does not prescribe a particular attack lifecycle methodology but provides sufficient conceptual flexibility to represent attack progression regardless of the attack framework adopted by the implementing platform. This abstraction allows the SKM to remain compatible with recognised cybersecurity knowledge bases while avoiding unnecessary dependencies on individual standards. A central characteristic of this domain is the representation of adversarial behaviour. Behaviour is understood as the observable operational pattern through which threat actors employ techniques, malware, offensive tools, exploits, scripts, infrastructure, and procedures to achieve specific objectives. Rather than describing these elements independently, the semantic model captures their relationships as components of coordinated attack behaviour. Attack techniques may therefore reuse common tools, exploit multiple vulnerabilities, trigger different operational events, or appear across unrelated campaigns while maintaining their own semantic identity within the model. This behavioural perspective enables CyberAid to analyse attack evolution, identify recurring adversarial patterns, and support behavioural correlation across heterogeneous intelligence sources.

The semantic organisation further distinguishes between attack techniques, attack procedures, and attack capabilities. Techniques describe the generic methods through which malicious objectives are achieved, whereas procedures represent their concrete implementation by specific adversaries within individual campaigns or operational contexts. Attack capabilities describe the resources, sophistication, infrastructure, or specialised competencies required to perform particular malicious activities. Maintaining this semantic distinction allows the SKM to represent both generic cybersecurity knowledge and adversary-specific operational behaviour without introducing conceptual ambiguity. It also facilitates behavioural comparison between campaigns, enabling common techniques to be recognised even when implemented differently by unrelated threat actors. Another important objective of the Attack & Adversarial Behaviour Domain is to represent offensive artefacts involved in attack execution. Malware families, offensive

software, exploitation frameworks, command-and-control infrastructures, scripts, payloads, and other attack-enabling resources constitute semantic entities that contribute to the operational description of malicious behaviour. These entities are not modelled solely as technical artefacts but as components participating in broader attack operations. Consequently, their semantic relationships extend beyond simple associations with malware samples or exploit code, allowing CyberAId to represent how offensive resources are reused across campaigns, adapted to different operational contexts, or combined with alternative attack techniques. The semantic model also supports the representation of attack evolution over time. Adversarial behaviour is inherently adaptive, with attackers continuously modifying techniques, exploiting newly discovered vulnerabilities, adopting emerging offensive tools, or changing operational objectives in response to defensive measures. The SKM therefore preserves temporal information describing the evolution of attack behaviour while maintaining historical relationships between successive observations. This capability supports longitudinal threat analysis, behavioural trend identification, and the continuous refinement of knowledge as additional intelligence becomes available. Within the CyberAId ecosystem, the Attack & Adversarial Behaviour Domain establishes the operational bridge between strategic threat intelligence and downstream cybersecurity domains. Attack behaviour provides the semantic context through which vulnerabilities become exploited, operational events are generated, incidents emerge, and defensive controls are activated. Consequently, this domain serves as the operational core of the SKM, connecting strategic adversarial knowledge with the technical evidence and defensive activities represented throughout the remaining semantic domains.

The proposed semantic specification has been designed to remain sufficiently generic to accommodate different sources of attack knowledge while preserving semantic consistency across the CyberAId platform. Established cybersecurity knowledge frameworks describing attack techniques, adversarial tactics, malware behaviour, exploitation patterns, or offensive methodologies may therefore be aligned with the SKM without requiring direct replication of their internal structures. Instead, these external knowledge sources provide complementary semantic information that enriches the common conceptual representation established by CyberAId, allowing the model to evolve alongside the cybersecurity landscape while maintaining implementation independence.

3.3.4 Vulnerability & Exposure Domain

The Vulnerability & Exposure Domain defines the semantic structures through which weaknesses, exploitable conditions, attack surfaces, and affected technological assets are represented within the CyberAId SKM. While the Threat Representation Domain identifies the strategic origin of malicious activities and the Attack & Adversarial Behaviour Domain models how adversaries conduct attacks, the Vulnerability & Exposure Domain explains where attacks can succeed by describing the technical conditions that enable malicious actions to compromise systems, services, applications, infrastructures, or digital assets. Consequently, this domain establishes the semantic bridge between adversarial behaviour and the operational environment in which attacks materialise. Within the proposed semantic architecture, vulnerabilities are represented as intrinsic weaknesses affecting software components, hardware platforms, communication infrastructures, cloud services, applications, configurations, identities, or operational processes. A vulnerability is considered an inherent characteristic of an affected entity and exists independently of whether exploitation has occurred. This distinction allows the SKM to separate latent weaknesses from observed attack activities, preserving the conceptual independence between system deficiencies and adversarial behaviour while enabling explicit semantic relationships between the two domains. Closely associated with vulnerabilities is the concept of exposure, which represents the operational conditions that make a vulnerability accessible to potential attackers. Exposure extends beyond the existence of a technical weakness by describing the environmental, architectural, configuration, accessibility, or operational circumstances that increase the likelihood of successful exploitation. An exposed service, an externally reachable application programming interface, an improperly configured cloud resource, excessive user privileges, insecure network segmentation, or publicly accessible administrative interfaces constitute representative examples of exposure conditions. By modelling exposure independently from vulnerabilities, the SKM provides a richer and more realistic representation of attack opportunities, supporting contextual reasoning that extends beyond software defects alone.

The proposed semantic specification therefore distinguishes between structural weaknesses and operational exposure. Structural weaknesses describe flaws embedded within systems or services, whereas exposure captures the operational context that determines whether those weaknesses can realistically be exploited. This distinction enables CyberAId to represent situations in which identical vulnerabilities may produce different cybersecurity implications depending on deployment architecture, business criticality, network topology, identity management policies, security controls, or operational configurations. Consequently, vulnerability assessment becomes intrinsically contextual rather than relying exclusively on technical severity metrics. The Vulnerability & Exposure Domain also introduces the concept of the attack surface as a semantic abstraction describing the collection of technological components, interfaces, services, communication channels, identities, APIs, cloud resources, and externally accessible assets that may become targets of adversarial activity. Rather than modelling attack surfaces as static inventories, the SKM represents them as dynamic semantic entities capable of evolving as infrastructures, configurations, business services, or operational environments change over time. This dynamic representation enables CyberAId to preserve an accurate view of organisational exposure while supporting continuous cybersecurity monitoring and knowledge evolution.

A further objective of this domain is the semantic representation of affected technological assets. Vulnerabilities and exposure conditions rarely exist independently of the systems they affect. Consequently, every weakness is expected to maintain explicit semantic relationships with the corresponding software components, operating systems, virtual infrastructures, cloud environments, financial applications, communication services, identity infrastructures, or operational technologies that constitute the organisational attack surface. This association allows the SKM to contextualise technical weaknesses according to the operational significance of the affected assets rather than considering vulnerabilities as isolated technical findings. The Vulnerability & Exposure Domain has also been designed to support lifecycle-oriented knowledge management. Vulnerabilities evolve continuously as new weaknesses are discovered, patches become available, exploit code is published, security advisories are updated, and mitigation measures are implemented. Similarly, exposure conditions may change due to infrastructure modifications, architectural redesign, cloud migrations, security hardening activities, or evolving operational requirements. The semantic model therefore preserves temporal information describing the discovery, publication, modification, remediation, verification, and retirement of vulnerabilities and exposure conditions, enabling CyberAId to analyse cybersecurity posture as an evolving operational state rather than a static snapshot.

Since CyberAId focuses on the financial sector, the semantic representation of vulnerabilities extends beyond conventional IT infrastructures. The proposed specification accommodates vulnerabilities affecting payment infrastructures, online banking platforms, financial APIs, digital identity services, customer authentication mechanisms, transaction processing systems, mobile banking applications, cloud-native financial services, third-party service providers, and interconnected financial ecosystems. This financial contextualisation enables vulnerabilities to be interpreted according to their operational importance within banking and financial environments while supporting downstream risk analysis and regulatory compliance activities. The Vulnerability & Exposure Domain also provides the semantic foundation upon which subsequent cybersecurity activities are interpreted. Attack techniques exploit vulnerabilities, operational events reveal exploitation attempts, incidents emerge from successful compromises, and defensive controls aim to reduce exposure or eliminate exploitable weaknesses. Consequently, this domain constitutes the technical pivot between offensive behaviour and defensive cybersecurity operations, enabling CyberAId to correlate threat intelligence, vulnerability knowledge, operational observations, and mitigation activities through a common semantic representation.

Eventually, the proposed specification maintains implementation independence while remaining compatible with recognised vulnerability management practices and cybersecurity knowledge ecosystems. Public vulnerability catalogues, weakness taxonomies, software identification schemes, exposure assessment methodologies, asset inventories, and security configuration repositories may therefore contribute complementary knowledge to the SKM without constraining its conceptual organisation. Rather than replicating existing standards, the Vulnerability & Exposure Domain provides the semantic abstraction

through which heterogeneous vulnerability information can be represented consistently, contextualised according to organisational environments, and interconnected with the remaining semantic domains of the CyberAid ecosystem.

3.3.5 Detection, Observation & Incident Domain

The Detection, Observation & Incident Domain defines the semantic structures through which cybersecurity observations, monitoring outputs, detection artefacts, investigative findings, and operational incidents are represented within the CyberAid SKM. Whereas the preceding semantic domains describe the strategic origin of cyber threats, the operational execution of attacks, and the technical conditions that enable successful exploitation, this domain captures how malicious activities become observable, how they are interpreted, and how they evolve into operational cybersecurity incidents. Consequently, it provides the semantic representation of the continuous monitoring and situational awareness capabilities required to transform heterogeneous cybersecurity evidence into actionable operational knowledge. The proposed specification distinguishes three complementary semantic layers within the operational detection process. The first layer concerns observations, representing raw evidence collected from cybersecurity sensors, monitoring systems, threat intelligence feeds, security logs, endpoint agents, network infrastructures, cloud services, identity providers, or external information sources. Observations constitute factual descriptions of events or conditions without necessarily implying malicious intent. Authentication attempts, process executions, configuration changes, network connections, API requests, file modifications, telemetry measurements, or newly published threat intelligence reports all represent observations that may contribute to the operational understanding of the cybersecurity environment. By modelling observations independently from analytical conclusions, the SKM preserves objective evidence while enabling subsequent interpretation through additional semantic relationships.

The second semantic layer corresponds to detection, where individual observations are analysed, correlated, enriched, and interpreted to identify potentially malicious behaviour. Detection represents the analytical process through which multiple observations are transformed into security-relevant knowledge by applying detection logic, behavioural correlation, anomaly identification, intelligence enrichment, contextual reasoning, or predefined analytical rules. A detection therefore constitutes a semantic conclusion supported by one or more observations rather than a direct representation of individual technical events. This distinction allows CyberAid to preserve analytical traceability by explicitly linking every detection to the underlying evidence from which it originates. The highest operational abstraction within this domain is represented by the concept of an incident. Incidents correspond to confirmed cybersecurity situations that require investigation, response, containment, recovery, or organisational decision-making. Unlike observations or detections, incidents integrate technical evidence with operational context, business impact, affected assets, organisational responsibilities, response activities, and post-incident knowledge. The SKM therefore represents incidents as evolving semantic entities that progressively accumulate evidence, analytical findings, remediation actions, and organisational knowledge throughout the complete incident response lifecycle. In addition to observations, detections, and incidents, the SKM may represent complementary operational artefacts such as findings, alerts, and notifications through common semantic definitions. A finding represents a structured, evidence-supported analytical result produced during cybersecurity analysis or investigation. An alert represents a prioritised operational indication derived from one or more observations, detections, or findings that may require analyst attention or subsequent action. A notification represents the communication of relevant cybersecurity, incident, or compliance information to an authorised human or system consumer according to the applicable operational and governance context. These concepts are defined at the semantic level to ensure consistent interpretation across CyberAid components; implementation-specific payload schemas, trigger logic, transport mechanisms, and notification workflows remain responsibilities of the corresponding integration and operational components. An essential objective of this domain is the representation of evidence provenance. Cybersecurity decisions depend heavily on the reliability, origin, integrity, and temporal validity of collected evidence. Consequently, observations preserve explicit relationships with their originating systems, collection mechanisms, timestamps, confidence assessments, processing history, and applicable information handling constraints. This provenance-aware representation enables CyberAid to

maintain complete traceability between raw observations, analytical detections, investigative findings, and final incident assessments, thereby supporting explainability, auditability, and evidence-based cybersecurity operations. Beyond representing operational events, the Detection, Observation & Incident Domain also models investigative knowledge generated throughout cybersecurity operations. During incident investigations, analysts continuously produce hypotheses, intermediate findings, evidence correlations, forensic artefacts, impact assessments, containment decisions, and remediation recommendations. These analytical products represent valuable organisational knowledge that extends beyond the incident itself. The SKM therefore preserves investigative artefacts as first-class semantic entities, enabling future investigations to reuse historical knowledge, identify recurring attack patterns, and support continuous learning across cybersecurity operations.

The proposed semantic architecture further accommodates heterogeneous detection methodologies. Cybersecurity observations may originate from signature-based monitoring, behavioural analytics, anomaly detection, rule-based correlation, graph analytics, machine learning models, cyber threat intelligence enrichment, or AI-assisted reasoning mechanisms. Rather than prescribing a particular analytical methodology, the SKM focuses on representing the semantic relationships between observations, detections, and incidents independently of the analytical techniques through which they were generated. This technology-independent approach allows CyberAid to integrate heterogeneous monitoring capabilities while maintaining a common conceptual representation of operational cybersecurity knowledge. Particular consideration is also given to the financial sector, where operational incidents frequently extend beyond purely technical consequences. Security observations may involve fraudulent transactions, compromised payment systems, suspicious authentication behaviour, abnormal customer activity, insider actions, or disruptions affecting critical financial services. Consequently, operational incidents are semantically associated not only with technical assets but also with business services, financial processes, regulatory obligations, customer-facing operations, and organisational resilience objectives. This multidimensional representation enables CyberAid to assess cybersecurity events according to both their technical characteristics and their operational significance within financial institutions. The Detection, Observation & Incident Domain also plays a central role in the continuous evolution of the SKM. Operational observations continuously enrich the knowledge ecosystem by confirming threat intelligence, identifying previously unknown attack techniques, validating exposure conditions, refining detection logic, or revealing emerging adversarial behaviour. Every confirmed incident therefore contributes additional semantic knowledge that may influence the Threat Representation Domain, the Attack & Adversarial Behaviour Domain, the Vulnerability & Exposure Domain, and the Defensive Knowledge Domain. Through this feedback mechanism, the SKM evolves dynamically as operational evidence accumulates, supporting continuous knowledge refinement and adaptive cybersecurity intelligence.

Thus, the proposed specification establishes this domain as the operational convergence point of the CyberAid semantic architecture. Threat intelligence, attack behaviour, vulnerabilities, defensive controls, financial assets, and organisational context all become interconnected through operational observations and incident knowledge, enabling CyberAid to transform heterogeneous cybersecurity information into actionable intelligence suitable for graph analytics, AI-assisted reasoning, decision support, compliance assessment, and coordinated cyber defence.

3.3.6 Defensive Knowledge Domain

The Defensive Knowledge Domain defines the semantic structures through which cybersecurity protection mechanisms, defensive capabilities, response procedures, recovery activities, and organisational resilience measures are represented within the CyberAid SKM. While the previous semantic domains progressively describe adversaries, attack execution, technical exposure, and operational observations, this domain captures how organisations prevent, detect, contain, respond to, recover from, and continuously improve against cyber threats. Consequently, it represents the defensive dimension of the SKM, enabling CyberAid to model cybersecurity not only from the perspective of malicious activities but also through the capabilities employed to protect critical financial infrastructures and services.

Within the proposed semantic architecture, defensive knowledge extends beyond the traditional representation of security controls. Modern cyber defence relies on a combination of preventive, detective, corrective, compensating, recovery, governance, and adaptive mechanisms that operate continuously throughout the cybersecurity lifecycle. The SKM therefore represents defence as a dynamic and evolving capability composed of multiple complementary semantic entities that collectively contribute to organisational cyber resilience. This approach allows CyberAId to describe not only individual security mechanisms but also the relationships between defensive capabilities, operational procedures, and the cybersecurity knowledge they consume or produce. The specification distinguishes several complementary categories of defensive knowledge. Preventive capabilities reduce the likelihood of successful attacks by strengthening systems, infrastructures, identities, applications, and operational processes before malicious activities occur. These capabilities include technical safeguards, secure configurations, authentication mechanisms, access control policies, encryption services, network segmentation strategies, software hardening measures, vulnerability remediation processes, and secure development practices. Their semantic representation enables the SKM to associate preventive controls directly with the vulnerabilities, exposure conditions, assets, or business services they protect, thereby supporting traceable and explainable defence planning.

A second category encompasses detective capabilities, whose purpose is to identify suspicious or malicious behaviour as early as possible. Rather than modelling detections themselves—which are represented within the Detection, Observation & Incident Domain—the Defensive Knowledge Domain represents the mechanisms responsible for producing those detections. Detection rules, behavioural models, anomaly detection algorithms, correlation logic, monitoring policies, intelligence enrichment procedures, and automated analytical workflows are therefore modelled as defensive assets that continuously analyse operational observations and generate security-relevant findings. This distinction preserves a clear separation between operational evidence and the analytical capabilities responsible for interpreting that evidence. The semantic model further introduces response and containment capabilities, representing the operational actions executed after malicious activities have been detected. Incident response procedures, containment actions, forensic acquisition processes, system isolation mechanisms, identity revocation, credential rotation, malware eradication, communication procedures, and recovery workflows are represented as independent semantic entities capable of interacting with incidents, assets, vulnerabilities, and organisational responsibilities. Rather than viewing response as a single activity, the SKM represents it as an orchestrated sequence of coordinated actions whose objective is to minimise operational disruption while preserving evidence and restoring organisational functionality.

An important concept introduced by this domain is the notion of defensive orchestration. Modern cybersecurity operations increasingly rely on coordinated execution of multiple defensive actions through predefined or adaptive workflows. Consequently, the SKM introduces semantic entities capable of representing playbooks, response workflows, automation strategies, orchestration policies, decision points, dependencies, and execution conditions. These entities enable CyberAId to represent defensive knowledge not merely as static documentation but as executable organisational knowledge capable of supporting automated or semi-automated cybersecurity operations. This representation is particularly relevant for environments integrating Security Orchestration, Automation and Response (SOAR) capabilities, AI-assisted decision support, or intelligent cyber defence agents. Beyond operational response, the Defensive Knowledge Domain also captures organisational learning and continuous improvement. Every cybersecurity incident provides valuable knowledge regarding the effectiveness of existing defensive capabilities, the adequacy of security controls, emerging attack techniques, and organisational preparedness. The SKM therefore preserves defensive lessons learned, recommended improvements, capability assessments, security maturity observations, validation outcomes, and knowledge derived from post-incident reviews as explicit semantic entities. By representing these elements within the knowledge ecosystem, CyberAId enables defensive knowledge to evolve continuously alongside the threat landscape rather than remaining static after incident resolution.

Given the regulatory environment of the financial sector, the proposed semantic model also represents governance-oriented defensive capabilities. Security policies, operational procedures, regulatory controls,

compliance mechanisms, audit requirements, resilience measures, business continuity plans, disaster recovery strategies, and organisational responsibilities constitute complementary defensive knowledge that supports both technical cybersecurity operations and regulatory compliance. Associating these governance-oriented entities with technical controls and operational incidents enables CyberAId to bridge cybersecurity operations with organisational governance and regulatory accountability without introducing conceptual fragmentation between technical and non-technical defence. The Defensive Knowledge Domain has also been designed to support adaptive cybersecurity. Defensive mechanisms are expected to evolve continuously according to newly acquired threat intelligence, vulnerability disclosures, attack observations, operational incidents, regulatory changes, and technological developments. Consequently, defensive capabilities preserve explicit semantic relationships describing their applicability, effectiveness, validation status, lifecycle evolution, implementation context, dependencies, and historical performance. This adaptive representation enables CyberAId to evaluate the effectiveness of defensive strategies over time while supporting evidence-based improvement of organisational cyber resilience.

Within the overall semantic architecture, the Defensive Knowledge Domain represents the final operational stage of the cybersecurity lifecycle while simultaneously acting as the starting point for future knowledge evolution. Defensive actions generate new operational observations, validate threat intelligence, refine attack knowledge, influence vulnerability management activities, improve detection capabilities, and enrich organisational experience. Through these feedback mechanisms, the SKM becomes a continuously evolving knowledge ecosystem in which defensive knowledge contributes not only to operational protection but also to the long-term maturity of the CyberAId platform. The proposed specification remains independent of specific cybersecurity frameworks, security products, orchestration technologies, or implementation methodologies. Security standards, control catalogues, incident response frameworks, governance models, and best-practice recommendations may all contribute complementary knowledge to the Defensive Knowledge Domain while preserving the technology-independent conceptual organisation defined by the SKM. This approach ensures that defensive knowledge can evolve together with emerging cybersecurity practices without requiring structural modifications to the underlying semantic architecture.

3.3.7 Cross-Domain Semantic Integration

The semantic domains introduced throughout this chapter collectively establish the conceptual foundation of the CyberAId SKM. While each domain represents a distinct perspective of cybersecurity knowledge, their true value emerges through their semantic integration into a unified knowledge ecosystem capable of representing complex cybersecurity situations from multiple complementary viewpoints. Rather than functioning as independent knowledge repositories, the semantic domains continuously exchange contextual information, progressively enriching one another and enabling CyberAId to construct comprehensive representations of threats, attacks, vulnerabilities, operational events, defensive activities, and business impacts. The SKM therefore adopts a cross-domain semantic integration approach, where cybersecurity knowledge is organised as an interconnected network of semantically related entities rather than a collection of isolated concepts. Every semantic entity contributes contextual information that extends beyond its own domain and supports the interpretation of entities belonging to other domains. Consequently, strategic threat intelligence can be associated with operational attack behaviour, attack techniques can be connected to exploited vulnerabilities, vulnerabilities can be linked to affected financial assets, operational incidents can reference defensive actions, and mitigation strategies can ultimately influence organisational resilience and future threat assessments. This interconnected organisation enables CyberAId to represent cybersecurity as an evolving operational ecosystem rather than a sequence of disconnected technical observations. A fundamental objective of cross-domain semantic integration is the preservation of context. Individual cybersecurity observations rarely provide sufficient information to support informed operational decisions. An isolated vulnerability, for example, does not necessarily indicate organisational exposure unless it is associated with reachable assets, relevant attack techniques, capable threat actors, operational detections, and corresponding defensive measures. Similarly, an observed malicious indicator gains significantly greater operational value when semantically connected with attack campaigns, affected financial services, previous incidents, existing mitigation controls, and historical intelligence. By preserving these contextual relationships, the SKM enables CyberAId to interpret

cybersecurity information according to its operational significance rather than its isolated technical characteristics.

Cross-domain integration also enables progressive knowledge enrichment. Cybersecurity knowledge evolves continuously as new intelligence reports, vulnerability disclosures, attack observations, investigation findings, regulatory updates, and defensive improvements become available. Instead of replacing existing information, newly acquired knowledge incrementally enriches the semantic relationships already established within the SKM. Threat actors become associated with newly identified campaigns, attack techniques become linked to additional malware families, vulnerabilities acquire updated exploitability information, incidents accumulate further evidence, and defensive capabilities evolve according to operational experience. This continuous enrichment process allows the knowledge ecosystem to mature organically while preserving historical context, provenance information, and semantic consistency across successive knowledge updates.

The proposed integration approach also supports semantic interoperability across heterogeneous cybersecurity information sources. Modern cybersecurity environments combine information originating from cyber threat intelligence platforms, vulnerability repositories, monitoring infrastructures, cloud services, identity management systems, regulatory frameworks, security operations centres, incident response platforms, and organisational knowledge repositories. These sources frequently employ different terminologies, data models, identifiers, and representation mechanisms. The SKM acts as a common semantic abstraction layer that harmonises these heterogeneous knowledge sources through shared conceptual definitions and explicit semantic relationships, enabling CyberAid to integrate diverse information without requiring complete structural uniformity between individual data providers. Another important characteristic of cross-domain integration concerns the support of knowledge correlation. The semantic architecture enables analytical processes to identify relationships that extend beyond direct technical dependencies. Multiple observations originating from different monitoring systems may collectively indicate the execution of a single attack campaign. Apparently unrelated vulnerabilities may affect the same critical financial service. Distinct incidents occurring across different institutions may reveal common adversarial infrastructure or coordinated threat activity. By establishing explicit semantic relationships between entities originating from different domains, CyberAid enables higher-level reasoning processes capable of discovering operational patterns that would remain hidden if each domain were analysed independently. The integrated semantic architecture further provides the conceptual foundation for graph-based knowledge representation. Although the SKM remains independent of specific implementation technologies, its interconnected structure naturally supports representation as a knowledge graph, where semantic entities correspond to nodes and their relationships define the edges connecting different knowledge domains. Such graph-oriented representations facilitate complex traversal operations, contextual exploration, dependency analysis, path discovery, and advanced correlation queries without constraining the conceptual organisation of the model to any particular graph database or ontology management technology. Consequently, graph technologies may serve as implementation options rather than architectural requirements.

Cross-domain semantic integration also establishes the necessary foundation for AI-assisted cybersecurity reasoning. Modern analytical capabilities increasingly rely on contextual knowledge rather than isolated cybersecurity observations. Retrieval-Augmented Generation (RAG), graph-enhanced language models, intelligent cyber agents, recommendation systems, decision-support mechanisms, and explainable AI algorithms all require structured contextual information capable of connecting strategic intelligence, technical evidence, operational knowledge, business assets, and defensive capabilities. By providing an integrated semantic representation of these complementary knowledge domains, the SKM enables AI components to retrieve coherent contextual information while preserving traceability, provenance, and explainability throughout the analytical process. Within the financial sector, cross-domain integration becomes particularly significant due to the complexity of modern financial ecosystems. Cybersecurity events rarely affect only technical infrastructures; they frequently influence payment services, customer authentication mechanisms, transactional integrity, operational continuity, regulatory compliance, reputational trust, and financial stability simultaneously. The integrated semantic architecture therefore

enables CyberAid to associate cybersecurity knowledge with financial assets, business processes, organisational responsibilities, and resilience objectives, allowing technical cybersecurity observations to be interpreted according to their broader operational consequences.

Finally, cross-domain semantic integration establishes the conceptual transition towards the subsequent sections of this deliverable. Having defined the principal semantic domains constituting the SKM, the following sections formally specify the semantic relationships, metadata structures, governance mechanisms, and implementation-oriented knowledge representations that operationalise the conceptual architecture described throughout this chapter. In this manner, the SKM progresses from a collection of independently defined semantic domains towards a fully interconnected and operational semantic ecosystem capable of supporting the broader CyberAid platform.

3.3.8 Core SKM class catalogue

Every catalogue entity inherits the common metadata, provenance, classification, confidence and lifecycle attributes specified in Section 3.6. The catalogue defines the minimum conceptual classes required by the initial SKM specification and may be extended through the controlled-extension process.

Table 3 - Class Catalogue

Entity	Core attributes	Requirement	Source domain
ThreatActor	element_id, name, description, aliases, confidence	element_id and name	Threat
Campaign	element_id, name, first_seen, last_seen, confidence	element_id and name	Threat
Indicator	element_id, pattern, pattern_type, valid_from, valid_until	element_id, pattern and pattern_type	CTI
AttackTechnique	element_id, name, external_identifier	element_id and name; external_identifier when mapped to ATT&CK	Adversarial behaviour
Vulnerability	element_id, name, external_identifier, severity, published	element_id and name; external_identifier when mapped to CVE or NVD	Vulnerability
Weakness	element_id, name, external_identifier	element_id and name; external_identifier when mapped to CWE	Vulnerability

Entity	Core attributes	Requirement	Source domain
Asset	element_id, name, asset_type, owner, criticality	element_id, name and asset_type	Organisational
Observation	element_id, observed_at, source_reference, content	element_id, observed_at, source_reference and content	Detection
Finding	element_id, created, summary, evidence_references, confidence	element_id, created, summary and evidence_references	Detection
Incident	element_id, name, status, first_seen, last_seen	element_id and status	Incident
SecurityControl	element_id, name, control_type, external_identifier	element_id and name	Defensive
Risk	element_id, name, likelihood, impact, risk_level	element_id and name	Risk
RegulatoryRequirement	element_id, framework, reference, obligation_text	element_id, framework, reference and obligation_text	Regulatory
SourceArtefact	element_id, title, source_uri, publisher, published	element_id and title	Provenance

3.4 Financial, Risk and Regulatory Concepts

The cybersecurity semantic architecture defined in the previous sections establishes a comprehensive representation of cyber threats, adversarial behaviour, technical exposure, operational observations, and defensive capabilities. While these domains provide the technical foundation required to describe cybersecurity events, they do not fully capture the organisational environment in which such events occur. CyberAId has been specifically designed to support the financial sector, where cybersecurity incidents extend beyond technical infrastructures and directly affect critical business operations, digital financial services, customer trust, regulatory compliance, and operational resilience. Consequently, the SKM extends its semantic architecture by incorporating financial, organisational, risk, and regulatory knowledge into the same interconnected semantic ecosystem.

Rather than introducing an independent business ontology, the proposed specification enriches the existing cybersecurity semantic architecture with complementary domains that provide organisational and

operational context to technical cybersecurity knowledge. This contextualisation enables the SKM to associate cyber threats, vulnerabilities, incidents, and defensive activities with the financial assets they affect, the business processes they support, the organisational risks they introduce, and the regulatory obligations they influence. Through this multidimensional representation, cybersecurity information evolves from isolated technical observations into actionable organisational knowledge capable of supporting operational decision-making, cyber resilience, and compliance activities. The financial, risk, and regulatory concepts introduced in this section therefore constitute the business-oriented semantic layer of the CyberAid SKM. Together with the cybersecurity semantic domains defined in Section 3.3, they establish a unified semantic architecture capable of representing both the technical characteristics of cybersecurity events and their organisational consequences within financial institutions. This integrated modelling approach enables CyberAid to support contextual cyber-risk analysis, AI-assisted reasoning, graph-based knowledge exploration, regulatory traceability, and resilience-oriented decision support while preserving semantic consistency across the entire knowledge ecosystem.

The following subsections progressively define the principal semantic domains representing financial infrastructures, organisational assets, operational risk, regulatory obligations, and their semantic integration with the cybersecurity knowledge model.

3.4.1 Financial Domain

The Financial Domain introduces the semantic concepts required to represent the operational environment in which cybersecurity events occur within financial institutions. While traditional cybersecurity knowledge models primarily focus on technological infrastructures, CyberAid recognises that cyber threats ultimately target business operations, financial services, customer transactions, payment ecosystems, and digital banking infrastructures rather than technical assets alone. The purpose of this semantic domain is therefore to establish a structured representation of the financial ecosystem, enabling cybersecurity knowledge to be interpreted according to the operational realities of financial organisations. Within the SKM, financial entities are represented as organisational concepts describing the business services, financial products, digital infrastructures, and operational functions that collectively enable financial institutions to deliver secure and resilient services. Rather than modelling financial systems purely from an information technology perspective, the proposed specification captures the operational dependencies that exist between financial assets, customer-facing services, payment infrastructures, authentication mechanisms, settlement platforms, transaction processing environments, and supporting digital services. This representation enables cybersecurity observations to be associated with the operational processes they influence rather than solely with the technological components involved.

A fundamental modelling principle of the Financial Domain is the distinction between business assets and technical assets. Technical assets, including servers, databases, APIs, cloud infrastructures, applications, or communication networks, are represented within the cybersecurity domains as technological entities susceptible to compromise. The Financial Domain, however, represents the organisational functions these technical assets enable, such as payment processing, customer authentication, account management, digital identity verification, fraud detection, treasury operations, securities trading, interbank communications, and regulatory reporting. This distinction allows CyberAid to evaluate cybersecurity events according to their operational significance rather than limiting analysis to technical system failures. The semantic model also represents financial services as first-class organisational entities. Digital banking services, mobile banking applications, payment gateways, open banking interfaces, electronic wallets, instant payment mechanisms, card processing infrastructures, customer relationship services, and financial data exchange platforms are modelled independently from the technological infrastructures supporting them. Such separation enables CyberAid to distinguish disruptions affecting business functionality from failures occurring within individual technological components, thereby facilitating more accurate impact analysis and operational resilience assessments.

Another important characteristic of this domain concerns the representation of business processes. Financial institutions operate through highly interconnected operational workflows that involve customer onboarding, payment authorisation, transaction validation, settlement processing, identity verification,

anti-money laundering controls, fraud prevention, regulatory reporting, and business continuity procedures. These business processes establish semantic relationships between organisational responsibilities, technological infrastructures, financial services, and cybersecurity controls. Consequently, attacks targeting a single technical component may be analysed according to their potential impact on multiple business processes, enabling CyberAid to preserve complete operational context throughout the knowledge ecosystem. The Financial Domain further introduces semantic entities representing organisational stakeholders and service dependencies. Internal organisational units, operational teams, security operations centres, compliance functions, executive management, external service providers, cloud operators, payment processors, and third-party financial partners collectively contribute to the delivery of secure financial services. Representing these organisational dependencies enables the SKM to associate cybersecurity events with operational responsibilities, supporting coordinated incident response, resilience planning, and governance-oriented decision making. Because financial organisations increasingly operate within interconnected digital ecosystems, the proposed specification also represents relationships between institutions, external infrastructures, and shared financial services. Payment networks, clearing houses, financial messaging infrastructures, open banking ecosystems, identity federations, cloud service providers, outsourced technology partners, and cross-border financial services may therefore be represented as interconnected semantic entities rather than isolated organisational components. This ecosystem-oriented representation reflects the growing importance of supply-chain security, third-party risk, and operational resilience within modern financial environments. The Financial Domain has been intentionally designed to remain independent from specific banking standards, payment protocols, or financial information models. Existing financial ontologies, business process catalogues, digital banking architectures, and financial messaging standards may contribute complementary semantic information without constraining the conceptual organisation defined by the SKM. This technology- and standard-independent approach allows CyberAid to accommodate diverse financial institutions, operational models, and evolving financial technologies while maintaining a stable semantic foundation capable of supporting future extensions of the platform.

Ultimately, the Financial Domain provides the organisational context that transforms generic cybersecurity knowledge into financial-sector cybersecurity intelligence. By explicitly associating cyber threats, vulnerabilities, operational incidents, and defensive capabilities with business services, financial assets, operational processes, and organisational responsibilities, the SKM enables CyberAid to reason about cybersecurity according to business impact, operational resilience, and service continuity rather than technical compromise alone. This semantic contextualisation establishes the basis upon which the subsequent Risk Representation Domain can evaluate the consequences of cybersecurity events from both operational and organisational perspectives.

3.4.2 Organisational Assets & Business Processes Domain

The Organisational Assets & Business Processes Domain extends the business-oriented semantic representation of the CyberAid SKM by introducing the organisational entities, operational resources, and business activities that collectively define the operational environment of financial institutions. While the Financial Domain represents the services and infrastructures delivered by financial organisations, this semantic domain focuses on the internal organisational elements that enable those services to operate. Its purpose is to establish an explicit semantic representation of the assets, organisational structures, operational processes, responsibilities, and dependencies that determine the operational resilience of financial institutions and provide the business context required for cybersecurity decision-making. Within the SKM, an organisational asset is defined as any resource whose confidentiality, integrity, availability, authenticity, or operational continuity contributes to the successful execution of organisational objectives. Unlike purely technical asset inventories, organisational assets encompass technological infrastructures, business information, operational capabilities, digital services, human resources, organisational knowledge, communication channels, third-party services, and supporting operational functions. This broader interpretation enables CyberAid to evaluate cybersecurity events according to the organisational value of affected resources rather than limiting analysis to technical infrastructure alone.

The proposed semantic specification introduces a clear distinction between technical assets, information assets, and business assets. Technical assets include physical and virtual infrastructures, applications, cloud resources, databases, communication networks, identity management systems, APIs, and software platforms supporting organisational operations. Information assets represent structured and unstructured information whose protection is essential for financial operations, including customer information, transaction records, authentication credentials, financial reports, audit evidence, operational documentation, and regulatory artefacts. Business assets represent the organisational capabilities and operational functions that enable institutions to deliver financial services, including customer onboarding, payment processing, fraud prevention, treasury management, regulatory reporting, customer support, and digital banking operations. Maintaining explicit semantic distinctions between these categories enables CyberAid to analyse cybersecurity events according to both technical impact and organisational consequence. Business processes constitute another fundamental component of this semantic domain. Financial institutions operate through highly interconnected workflows involving multiple organisational units, information assets, technological infrastructures, external partners, and regulatory obligations. The SKM therefore represents business processes as first-class semantic entities capable of describing operational activities, decision points, process dependencies, information exchanges, organisational responsibilities, and supporting digital services. Rather than viewing cybersecurity incidents as isolated technical failures, this representation enables CyberAid to evaluate how disruptions propagate across organisational operations and affect the continuity of critical financial services.

The semantic model also captures organisational dependencies that influence operational resilience. Critical business processes frequently depend upon shared infrastructures, third-party technology providers, outsourced operational services, cloud platforms, identity providers, payment gateways, telecommunications services, and external regulatory infrastructures. These dependencies introduce additional complexity into cybersecurity analysis, as attacks targeting external entities may indirectly affect the operational continuity of financial organisations. By explicitly modelling dependency relationships, the SKM enables CyberAid to represent cascading operational effects, supply-chain dependencies, and interconnected service ecosystems within a unified semantic framework. A further objective of this domain is the representation of criticality and business importance. Not all organisational assets contribute equally to institutional operations. Consequently, the SKM associates organisational assets and business processes with contextual attributes describing criticality, operational importance, business value, service priority, availability requirements, recovery objectives, ownership, and organisational responsibility. These semantic characteristics enable CyberAid to distinguish between assets whose compromise would produce limited operational consequences and those supporting essential financial functions whose disruption may significantly affect institutional resilience or customer trust.

The Organisational Assets & Business Processes Domain also supports the semantic representation of ownership and accountability. Every organisational asset and business process is associated with clearly defined organisational responsibilities, including asset owners, process owners, operational managers, security teams, compliance functions, executive stakeholders, and external service providers where applicable. This representation enables cybersecurity events to be linked directly with responsible organisational entities, facilitating coordinated incident management, governance processes, accountability, and evidence-based decision-making throughout the cybersecurity lifecycle. Because CyberAid addresses resilience within complex financial ecosystems, this semantic domain also captures inter-process relationships and operational dependencies. Business processes rarely operate independently; instead, they exchange information, invoke shared services, consume common technological infrastructures, and collectively support organisational objectives. A disruption affecting one operational process may therefore influence multiple dependent services, customer operations, or regulatory activities. The semantic representation of these relationships enables CyberAid to analyse the operational propagation of cybersecurity incidents and evaluate cascading impacts throughout interconnected organisational environments.

An additional characteristic of this domain concerns the representation of operational resilience objectives. Organisational assets and business processes are semantically associated with continuity requirements,

resilience objectives, recovery priorities, service-level expectations, and acceptable operational tolerances. Rather than representing resilience solely as a regulatory concept, the SKM integrates resilience directly into the semantic description of organisational operations, allowing cybersecurity observations and incidents to be evaluated according to their influence on business continuity and service availability. This approach aligns naturally with the resilience-oriented objectives of the financial sector while supporting future integration with risk and regulatory reasoning mechanisms. The Organisational Assets & Business Processes Domain therefore provides the operational context required to interpret cybersecurity knowledge according to organisational priorities rather than technical observations alone. By explicitly modelling assets, processes, dependencies, ownership, criticality, and resilience objectives, the SKM enables CyberAid to understand not only how cyber-attacks occur but also why they matter to the organisation, which operational functions they threaten, and which business capabilities require prioritised protection.

3.4.3 Risk Representation Domain

The Risk Representation Domain establishes the semantic structures through which cybersecurity knowledge is transformed into organisational risk knowledge within the CyberAid SKM. Building upon the cybersecurity, financial, and organisational domains introduced in the preceding sections, this semantic domain enables CyberAid to represent the potential consequences of cybersecurity events from both technical and business perspectives. Rather than considering risk as an isolated entity, the proposed specification models cyber risk as the outcome of semantic relationships between threats, adversarial behaviour, vulnerabilities, organisational assets, business processes, operational dependencies, defensive capabilities, and contextual information. This integrated representation allows cybersecurity observations to be progressively interpreted according to their potential impact on financial institutions and their operational resilience.

Unlike conventional risk models that frequently represent risk as a static combination of likelihood and impact, the CyberAid SKM adopts a knowledge-centric representation of risk. Risk is not treated as an independent object but as an emergent semantic construct that derives from the interaction of multiple knowledge domains. Threat actors create opportunities for malicious activity, attack techniques exploit vulnerabilities, vulnerabilities expose organisational assets, compromised assets disrupt business processes, and operational disruptions ultimately influence organisational objectives, customer services, regulatory obligations, and financial resilience. The Risk Representation Domain therefore captures these dependencies explicitly, enabling CyberAid to reason about risk through contextual knowledge rather than isolated numerical values. The proposed semantic model distinguishes between technical risk, operational risk, and business risk, recognising that cybersecurity events frequently evolve across multiple organisational layers. Technical risk reflects the possibility of technological compromise resulting from exploitable weaknesses or malicious activities. Operational risk captures the disruption of organisational processes, critical services, operational workflows, and supporting infrastructures. Business risk represents the broader organisational consequences that may arise from operational disruption, including financial loss, reputational damage, customer impact, legal exposure, regulatory penalties, and strategic implications. Maintaining explicit semantic distinctions between these categories enables the SKM to represent the progression of cyber risk from technical events towards organisational consequences while preserving traceability across every stage of the cybersecurity lifecycle.

A central objective of the Risk Representation Domain is the semantic modelling of risk scenarios. Rather than associating risk directly with isolated vulnerabilities or incidents, CyberAid represents risk scenarios as contextual descriptions of potential or realised cybersecurity situations affecting organisational operations. A risk scenario may therefore combine multiple threat actors, attack techniques, vulnerabilities, exposed assets, business services, operational dependencies, existing security controls, and external environmental conditions into a coherent semantic representation describing how a cybersecurity event may unfold within a financial institution. This scenario-oriented representation enables more realistic reasoning by preserving the contextual relationships that influence organisational exposure and resilience. The Risk Representation Domain further incorporates the concept of business impact as a semantic entity distinct from technical compromise. Business impact describes the operational and strategic consequences resulting from successful cyber-attacks or significant cybersecurity incidents. Such impacts may include

interruption of payment services, degradation of digital banking platforms, unauthorised financial transactions, loss of customer trust, disruption of trading activities, unavailability of critical services, operational delays, contractual breaches, financial losses, reputational harm, or regulatory reporting obligations. Representing these consequences explicitly enables CyberAId to prioritise cybersecurity activities according to organisational objectives rather than technical severity alone. Another important characteristic of this semantic domain concerns the representation of risk propagation. Modern financial ecosystems are composed of highly interconnected infrastructures where cybersecurity incidents rarely remain confined to a single technological component. Compromise of one organisational asset may propagate across interconnected systems, shared infrastructures, outsourced services, business processes, customer-facing platforms, or external financial ecosystems. The SKM therefore represents dependency chains and propagation pathways as semantic relationships capable of describing how technical failures evolve into wider operational disruptions. This capability supports impact analysis, resilience assessment, and dependency-aware decision making across complex financial environments.

The proposed specification also introduces semantic entities representing risk treatment and decision support. Risk does not terminate at its identification but continues through assessment, prioritisation, acceptance, mitigation, transfer, monitoring, and periodic reassessment. Consequently, the SKM associates risk scenarios with treatment strategies, responsible organisational functions, mitigation plans, residual risk evaluations, review activities, and implementation status. Rather than embedding quantitative risk methodologies within the semantic model, the proposed specification provides the conceptual structures required to represent the complete organisational lifecycle of cyber risk while remaining independent of specific assessment methodologies or scoring techniques. An equally important objective of this domain is the representation of uncertainty and confidence. Cybersecurity risk assessments are inherently dynamic and frequently rely on incomplete, evolving, or partially validated information. Threat intelligence may remain uncertain, vulnerabilities may have unknown exploitability, business dependencies may evolve over time, and operational impacts may change according to organisational circumstances. The SKM therefore preserves contextual information describing confidence, evidence quality, assessment maturity, temporal validity, and analytical assumptions associated with every risk representation. This provenance-aware approach enables CyberAId to support explainable decision-making while avoiding deterministic interpretations of inherently uncertain cybersecurity knowledge. Because CyberAId targets highly regulated financial environments, the Risk Representation Domain also connects cybersecurity risk with operational resilience objectives. Risk scenarios are semantically associated with business continuity requirements, critical service availability, recovery objectives, resilience capabilities, and organisational tolerance for disruption. This contextualisation enables CyberAId to evaluate cyber risk not solely according to the probability of technical compromise but according to the institution's ability to maintain essential financial services under adverse conditions. Consequently, risk representation becomes directly aligned with resilience-oriented decision making and organisational preparedness rather than exclusively supporting vulnerability prioritisation.

Finally, the Risk Representation Domain serves as the semantic bridge between cybersecurity operations and organisational governance. By connecting technical cybersecurity observations with business consequences, operational resilience, financial impact, and organisational decision-making, the SKM enables CyberAId to provide contextual knowledge suitable for AI-assisted cyber-risk analysis, graph-based reasoning, strategic planning, and executive decision support. Rather than replacing established risk management methodologies, the semantic model provides the common knowledge representation through which heterogeneous risk-related information can be interpreted consistently across the CyberAId ecosystem.

3.4.4 Regulatory & Compliance Domain

The Regulatory & Compliance Domain establishes the semantic structures through which legal, regulatory, governance, and compliance knowledge is incorporated into the CyberAId SKM. While the previous semantic domains describe cybersecurity, financial operations, organisational assets, and cyber risk, this domain introduces the regulatory context that governs how financial institutions are expected to manage cybersecurity, operational resilience, information security, and digital trust. Rather than representing

regulations as isolated documents or static reference material, the proposed specification models regulatory knowledge as an integral component of the semantic ecosystem, enabling CyberAid to associate cybersecurity activities directly with applicable legal obligations, governance requirements, and compliance evidence.

The Financial sector operates within one of the most demanding regulatory environments, where cybersecurity decisions must satisfy not only technical security objectives but also legal, supervisory, and operational resilience requirements. Consequently, the SKM represents regulatory knowledge as a collection of semantic entities describing obligations, requirements, control objectives, governance principles, reporting responsibilities, supervisory expectations, and evidence associated with cybersecurity management. This representation enables technical cybersecurity observations to be interpreted according to their regulatory significance, supporting explainable compliance reasoning and facilitating regulatory traceability throughout the cybersecurity lifecycle. A fundamental modelling principle of this domain is the explicit distinction between regulatory requirements and compliance status. Regulatory requirements represent the obligations imposed by applicable legal frameworks, standards, or supervisory authorities. Compliance status, on the other hand, describes the extent to which an organisation demonstrates conformity with those obligations based on available evidence, implemented controls, documented procedures, operational practices, or assessment activities. This separation enables CyberAid to represent compliance as a continuously evolving organisational condition rather than a binary attribute, allowing regulatory posture to change dynamically as cybersecurity knowledge, operational evidence, or organisational controls evolve.

The proposed semantic model introduces regulatory obligations as first-class semantic entities. These obligations may describe requirements related to ICT risk management, security governance, operational resilience, incident reporting, vulnerability management, access control, identity management, third-party risk, digital operational resilience testing, business continuity, disaster recovery, personal data protection, auditability, logging, evidence preservation, and organisational accountability. Rather than assigning these obligations directly to technical components, the SKM associates them with business processes, organisational assets, security controls, risk scenarios, operational incidents, and governance activities, thereby enabling comprehensive compliance contextualisation across multiple semantic domains. The Regulatory & Compliance Domain further introduces the concept of compliance evidence, recognising that regulatory conformity depends not only on implemented controls but also on the availability of verifiable evidence demonstrating their effectiveness. Policies, operational procedures, audit records, monitoring results, security logs, vulnerability assessments, penetration testing reports, incident records, risk assessments, resilience exercises, control validation activities, and management reviews therefore become semantic entities capable of supporting regulatory verification. By preserving explicit relationships between compliance evidence and the corresponding regulatory obligations, the SKM enables CyberAid to support traceability, audit preparation, and explainable compliance assessment.

An additional objective of this domain is the semantic representation of governance responsibilities. Cybersecurity governance within financial institutions extends beyond technical implementation and involves executive management, ICT functions, risk management teams, compliance officers, internal audit, business owners, operational managers, and external supervisory authorities. The SKM therefore captures organisational accountability through semantic relationships describing ownership, responsibility, approval, oversight, review, reporting, and decision-making activities. This governance-oriented representation allows CyberAid to associate cybersecurity events with organisational accountability while supporting management reporting and regulatory oversight.

The proposed specification also accommodates the coexistence of multiple regulatory frameworks within a common semantic architecture. Financial institutions are frequently required to satisfy overlapping requirements originating from European regulations, international standards, sector-specific guidance, and internal governance policies. Rather than modelling each framework independently, the SKM introduces a unified semantic representation capable of linking common cybersecurity concepts to multiple regulatory sources simultaneously. Consequently, a single security control or operational activity may demonstrate

compliance with several regulatory obligations without introducing duplicated semantic structures, supporting regulatory harmonisation and reducing conceptual fragmentation. The regulatory and standards baseline considered by CyberAId is documented in D2.1, where particular emphasis is placed on regulations and standards that are directly relevant to the objectives of CyberAId and the European financial ecosystem. The semantic model accommodates obligations derived from the Digital Operational Resilience Act (DORA), the General Data Protection Regulation (GDPR), the NIS2 Directive, the Cyber Resilience Act (CRA), the AI Act, as well as internationally recognised standards and frameworks such as ISO/IEC 27001, ISO/IEC 27002, ISO/IEC 27005, the NIST Cybersecurity Framework, and relevant supervisory guidance. These frameworks are not replicated within the SKM; instead, they provide authoritative reference sources whose requirements can be semantically linked to organisational assets, cybersecurity controls, risk scenarios, incidents, and compliance evidence. The Regulatory & Compliance Domain also supports continuous compliance monitoring. Organisational compliance is not considered a one-time assessment but an evolving operational condition that changes as infrastructures are modified, threats emerge, controls are updated, incidents occur, or regulations evolve. Consequently, compliance status is represented as a dynamic semantic construct capable of incorporating new evidence, updated assessments, regulatory revisions, control validations, and operational observations. This dynamic representation aligns naturally with CyberAId's objective of supporting continuous cyber resilience rather than periodic compliance verification.

Eventually, the Regulatory & Compliance Domain establishes the governance perspective of the SKM by connecting cybersecurity operations with legal obligations, organisational accountability, regulatory evidence, and supervisory expectations. By integrating regulatory knowledge into the broader semantic ecosystem, CyberAId enables cybersecurity information to be interpreted simultaneously from technical, operational, organisational, business, and regulatory viewpoints. This multidimensional representation supports explainable AI, graph-based reasoning, evidence-driven compliance assessment, and informed decision-making while preserving semantic consistency across the complete knowledge architecture.

3.4.5 Cross-Domain Business Integration

The semantic domains introduced throughout Section 3.4 collectively establish the business-oriented perspective of the CyberAId SKM by extending the technical cybersecurity architecture with financial, organisational, risk, and regulatory knowledge. Individually, these domains provide structured representations of the operational environment in which cybersecurity activities occur. Their full value, however, is realised through their semantic integration with the cybersecurity domains defined in Section 3.3, creating a unified knowledge ecosystem capable of representing cybersecurity from technical, organisational, operational, business, and governance perspectives simultaneously.

The SKM therefore adopts a cross-domain business integration approach, where cybersecurity knowledge is continuously enriched through business context rather than maintained separately from organisational information. Technical entities such as threats, attack techniques, vulnerabilities, security events, incidents, and defensive controls are not interpreted in isolation but are semantically associated with financial services, organisational assets, business processes, operational dependencies, risk scenarios, governance responsibilities, and regulatory obligations. This integrated representation enables CyberAId to interpret cybersecurity observations according to their operational significance and organisational consequences, transforming isolated technical knowledge into business-aware cyber intelligence. A fundamental objective of cross-domain business integration is the preservation of business context throughout the complete cybersecurity lifecycle. Technical cybersecurity observations rarely provide sufficient information for organisational decision-making unless they are interpreted according to the operational environment they affect. A vulnerability, for example, acquires substantially greater relevance when associated with a critical payment service, an essential customer authentication process, or an externally accessible financial platform. Similarly, a detected attack becomes considerably more meaningful when connected to affected business capabilities, operational dependencies, financial assets, organisational responsibilities, and applicable regulatory obligations. The SKM therefore preserves these semantic associations explicitly, enabling CyberAId to reason about cybersecurity events according to both technical characteristics and business impact. Cross-domain business integration also enables organisational knowledge correlation

across heterogeneous information domains. Cybersecurity operations continuously generate technical observations that may influence operational resilience, business continuity, regulatory compliance, and enterprise risk management. Likewise, organisational changes, new business services, regulatory updates, third-party dependencies, or modifications to operational processes may alter the cybersecurity posture of financial institutions. By integrating these complementary knowledge domains within a common semantic architecture, CyberAId supports bidirectional reasoning in which technical cybersecurity knowledge informs business decision-making while organisational context continuously refines cybersecurity analysis.

Another important characteristic of the proposed integration model concerns the representation of operational dependencies. Financial institutions operate through highly interconnected ecosystems involving customers, payment infrastructures, cloud providers, third-party technology suppliers, regulatory authorities, financial market infrastructures, outsourced operational services, and digital identity providers. Cybersecurity events affecting one component may therefore propagate across multiple organisational functions and external dependencies. The SKM captures these relationships through explicit semantic associations that preserve dependency chains, service interactions, organisational ownership, and business process connectivity. This representation enables CyberAId to support dependency-aware impact analysis, operational resilience assessment, and informed prioritisation of defensive activities. The integrated semantic architecture further establishes the foundation for business-aware cyber risk reasoning. Rather than evaluating cybersecurity exclusively according to technical severity, the SKM enables CyberAId to correlate threat intelligence, attack behaviour, technical exposure, business criticality, organisational dependencies, resilience objectives, and compliance obligations within a unified analytical framework. This contextual representation supports richer reasoning processes capable of identifying how technical events influence business operations, operational continuity, customer services, financial stability, and regulatory exposure. Consequently, cyber risk emerges as a contextual organisational property rather than a purely technical assessment.

Cross-domain business integration also enables regulatory contextualisation of cybersecurity knowledge. Organisational assets, business processes, security controls, incidents, and operational risks may all be associated with applicable regulatory requirements and corresponding compliance evidence. This multidimensional representation allows CyberAId to analyse cybersecurity events simultaneously from operational and regulatory perspectives, supporting governance activities, audit preparation, supervisory reporting, and continuous compliance monitoring without introducing independent compliance models disconnected from operational cybersecurity knowledge. The proposed integration approach has also been designed to support AI-assisted organisational reasoning. Modern AI systems require structured contextual information capable of connecting cybersecurity events with business semantics, organisational responsibilities, resilience objectives, and governance knowledge. By integrating these complementary domains into a unified semantic architecture, the SKM enables graph-based reasoning, Retrieval-Augmented Generation (RAG), intelligent cyber agents, explainable decision-support systems, and advanced knowledge exploration mechanisms to operate upon a common organisational representation rather than multiple disconnected information repositories. This integrated perspective significantly improves the quality, explainability, and operational relevance of AI-assisted cybersecurity analysis. An equally important characteristic of cross-domain business integration is its support for knowledge evolution. Financial organisations continuously adapt to changing business requirements, emerging technologies, evolving cyber threats, revised regulatory frameworks, and new operational dependencies. The SKM therefore treats business knowledge as an evolving semantic ecosystem rather than a static organisational description. Financial services, organisational structures, risk scenarios, governance obligations, resilience objectives, and operational relationships may all evolve over time while preserving semantic consistency, historical traceability, and contextual integrity across the knowledge ecosystem.

Ultimately, cross-domain business integration transforms the CyberAId SKM into a comprehensive representation of organisational cyber resilience. Rather than maintaining separate cybersecurity, business, risk, and compliance models, CyberAId establishes a single interconnected semantic architecture in which technical knowledge, organisational context, operational processes, governance responsibilities, and regulatory obligations coexist within the same knowledge ecosystem. This integrated approach

provides the conceptual foundation required for advanced analytics, explainable AI, graph-based reasoning, cyber-risk assessment, operational resilience evaluation, and evidence-driven decision support throughout the CyberAid platform.

3.5 Relationships and Semantic Associations

The CyberAid SKM is founded upon the principle that cybersecurity knowledge derives not only from the representation of individual concepts but also from the semantic relationships established between them. While the previous sections introduced the principal semantic domains and their associated entities, the present section specifies the relationship model that enables these entities to form an interconnected and semantically coherent knowledge ecosystem. Rather than considering semantic entities as isolated objects, the SKM represents them as components of an evolving network in which each entity acquires meaning through its explicit associations with other concepts, observations, organisational elements, and contextual information. Semantic relationships constitute the structural foundation of the SKM. They define how knowledge originating from different domains—including cyber threat intelligence, attack behaviour, vulnerabilities, organisational assets, business processes, operational incidents, defensive capabilities, risk scenarios, and regulatory obligations—is interconnected within a common conceptual architecture. Through these relationships, CyberAid is capable of representing complex cybersecurity situations as integrated knowledge structures instead of disconnected collections of technical information. The relationship model has been designed according to the principles of semantic consistency, interoperability, extensibility, explainability, and technology independence. Relationships are therefore defined at the conceptual level without prescribing implementation-specific graph schemas, ontology languages, database structures, or knowledge representation technologies. This abstraction enables the SKM to be implemented using different semantic technologies while preserving identical conceptual meaning across all CyberAid components. Rather than introducing an exhaustive catalogue of predefined predicates, the specification organises semantic relationships into categories according to their functional role within the knowledge ecosystem. This categorisation facilitates consistent modelling practices, supports interoperability between heterogeneous knowledge sources, and enables future extensions of the SKM without affecting its conceptual integrity. Furthermore, relationship semantics are intentionally separated from implementation details, allowing different storage technologies or reasoning mechanisms to realise the same conceptual relationships through alternative technical approaches.

The following subsections define the principal categories of semantic relationships adopted by the SKM, including structural, behavioural, contextual, temporal, and cross-domain associations. Together, these relationships establish the semantic foundation that enables CyberAid to correlate heterogeneous cybersecurity knowledge, preserve contextual information, support graph-based reasoning, facilitate AI-assisted knowledge exploration, and maintain semantic consistency throughout the evolution of the knowledge ecosystem.

3.5.1 Semantic Relationship Model

The CyberAid SKM is based on the principle that meaningful cybersecurity knowledge emerges through the explicit representation of relationships between semantic entities rather than through isolated descriptions of individual concepts. While semantic entities define the objects that compose the knowledge ecosystem, semantic relationships define how these entities interact, influence one another, evolve throughout the cybersecurity lifecycle, and collectively describe complex cybersecurity situations. Consequently, the Semantic Relationship Model constitutes the structural backbone of the SKM, enabling heterogeneous cybersecurity knowledge to be organised into a coherent, contextual, and machine-interpretable semantic architecture.

Within the proposed specification, relationships are considered first-class semantic constructs rather than simple links between entities. Every relationship expresses a well-defined semantic meaning describing the nature of the association between two or more concepts. These associations may represent structural dependencies, operational interactions, causal effects, temporal evolution, organisational context, regulatory applicability, or provenance information. By explicitly defining the semantics of each

relationship, the SKM preserves contextual meaning across different knowledge domains and supports consistent interpretation by both human analysts and automated reasoning mechanisms. The Semantic Relationship Model has been designed according to the principles of semantic consistency, technology independence, extensibility, interoperability, and explainability. Relationships are therefore defined conceptually without assuming a specific implementation technology, ontology language, graph database, or knowledge representation framework. Whether the SKM is implemented using ontology technologies, property graphs, semantic graphs, knowledge bases, or hybrid AI-oriented architectures, the conceptual meaning of each relationship remains unchanged. This separation between conceptual semantics and technical implementation ensures long-term maintainability of the model while facilitating interoperability between different CyberAId components. A key characteristic of the Semantic Relationship Model is the use of typed relationships. Rather than relying on generic associations between entities, every relationship expresses a specific semantic role that describes how the participating entities are connected. For example, a threat actor may execute an attack campaign, an attack technique may exploit a vulnerability, a vulnerability may affect an organisational asset, a security control may mitigate a vulnerability, or a regulatory requirement may apply to a business process. Typed relationships improve semantic precision, reduce ambiguity, and provide a richer foundation for graph analytics, knowledge correlation, and AI-assisted reasoning.

The proposed model further assumes that relationships are directional whenever semantic meaning depends upon the orientation of the association. Directionality enables the SKM to distinguish between conceptually different relationships that may otherwise appear identical. For example, an attack targets an organisational asset, whereas an asset does not target an attack. Likewise, a security control protects an asset, while the protected asset merely is protected by the corresponding control. Preserving directional semantics allows CyberAId to represent causal dependencies, operational workflows, and information flows with greater conceptual accuracy while supporting more expressive graph traversals and reasoning processes. In addition to directionality, the Semantic Relationship Model recognises that relationships may exhibit different cardinality patterns. Certain entities may participate in one-to-one associations, whereas others naturally establish one-to-many or many-to-many relationships. A single threat actor may conduct multiple campaigns, an attack campaign may employ several attack techniques, a vulnerability may affect numerous assets, and a single security control may mitigate multiple vulnerabilities across different organisational environments. Rather than constraining the model through rigid cardinality assumptions, the specification allows relationships to reflect the complexity of real-world cybersecurity ecosystems while maintaining semantic consistency across interconnected domains. The Semantic Relationship Model also supports relationship composition, whereby multiple individual relationships collectively describe higher-level cybersecurity situations. Rather than representing isolated associations, the SKM enables semantic chains linking strategic intelligence, adversarial behaviour, technical exposure, organisational context, operational incidents, defensive capabilities, business impact, and regulatory obligations into coherent knowledge structures. These relationship chains provide the contextual foundation for graph analytics, attack path reconstruction, cyber-risk reasoning, dependency analysis, and explainable AI throughout the CyberAId platform.

An equally important characteristic concerns the distinction between entity semantics and relationship semantics. Entities describe the characteristics of individual concepts, whereas relationships capture the interactions between them. This separation enables CyberAId to maintain a modular knowledge architecture in which new semantic entities, business concepts, intelligence sources, or regulatory frameworks may be incorporated without fundamentally altering the existing relationship model. Conversely, new relationship types may be introduced as cybersecurity knowledge evolves while preserving the integrity of previously defined semantic entities. The Semantic Relationship Model has also been designed to facilitate semantic interoperability across heterogeneous knowledge sources. Cyber threat intelligence platforms, vulnerability repositories, financial information systems, regulatory documents, monitoring infrastructures, and organisational knowledge bases frequently employ different internal data models and terminology. By defining relationships at the conceptual level, the SKM provides a common semantic layer capable of harmonising these heterogeneous sources without requiring identical

underlying data structures. This interoperability is essential for supporting integrated cybersecurity analysis across the diverse technological ecosystem envisioned by CyberAid.

Finally, the Semantic Relationship Model establishes the conceptual basis for the relationship categories introduced in the following subsections. Structural relationships define hierarchical and compositional associations between entities, behavioural relationships describe operational interactions and attack dynamics, contextual relationships capture organisational and business semantics, while temporal and provenance relationships preserve knowledge evolution, traceability, and lifecycle information. Together, these complementary relationship categories transform the SKM into an interconnected semantic ecosystem capable of supporting advanced cybersecurity analysis, graph-oriented knowledge exploration, AI-assisted reasoning, and continuous semantic evolution.

3.5.2 Structural Relationships

Structural relationships establish the semantic foundations through which entities are organised, classified, composed, and interconnected within the CyberAid SKM. Unlike behavioural or operational relationships, which describe interactions between entities during cybersecurity activities, structural relationships define the intrinsic organisation of knowledge by expressing hierarchical, compositional, dependency-oriented, and classification semantics. These relationships provide the conceptual framework that allows the SKM to organise heterogeneous cybersecurity, financial, organisational, and regulatory concepts into a coherent and navigable semantic architecture. The principal objective of structural relationships is to ensure that every semantic entity can be positioned consistently within the broader knowledge ecosystem. Rather than representing cybersecurity concepts as isolated objects, the SKM organises them into explicit semantic structures that capture inheritance, composition, ownership, dependency, and organisational hierarchy. Through these structural associations, individual concepts acquire additional contextual meaning by explicitly defining their position relative to other entities within the model. One of the most fundamental categories of structural relationships concerns hierarchical classification. Hierarchical relationships enable the SKM to organise concepts into taxonomies that describe generalisation and specialisation without introducing semantic ambiguity. Generic concepts may therefore be refined into progressively more specialised entities while preserving inheritance of their common semantic characteristics. For example, a generic security control may be specialised into preventive, detective, corrective, or compensating controls, while a general organisational asset may be refined into technical, information, or business assets. This hierarchical organisation facilitates semantic consistency, simplifies knowledge maintenance, and enables different levels of abstraction to coexist within the same knowledge ecosystem.

A second category comprises compositional relationships, which describe how complex semantic entities are constructed from simpler constituent components. Financial services, business processes, attack campaigns, organisational structures, and resilience capabilities frequently consist of multiple interconnected elements that collectively form a higher-level semantic entity. Compositional relationships enable the SKM to preserve these structures explicitly, allowing complex organisational concepts to be analysed both as complete entities and as collections of interrelated components. This capability is particularly important for representing financial ecosystems, where business services often depend upon numerous technological infrastructures, operational processes, external providers, and organisational functions. Structural relationships also capture ownership and responsibility structures. Organisational knowledge is characterised not only by technical dependencies but also by clearly defined accountability relationships describing ownership, stewardship, operational management, governance responsibilities, and administrative authority. Consequently, the SKM explicitly represents ownership relationships linking organisational assets, business processes, security controls, compliance obligations, and operational functions to the responsible organisational entities. These associations support governance, accountability, asset management, and operational coordination while providing contextual information that may subsequently be reused by risk assessment, incident response, and regulatory reasoning activities.

Another important category concerns dependency relationships, which describe structural reliance between entities without necessarily implying operational interaction. Business processes depend on supporting services, applications depend on underlying infrastructures, organisational capabilities rely on

specific assets, and financial services frequently require external providers or shared technological platforms. Dependency relationships enable the SKM to capture these architectural structures independently from the operational events that may occur during normal operation or cyber-attacks. This distinction allows CyberAId to analyse organisational architectures, identify potential single points of failure, and understand the structural context within which cybersecurity events occur. The proposed semantic model further supports the representation of containment and membership relationships. Many cybersecurity concepts naturally exist within larger organisational or operational contexts. Assets belong to organisational units, incidents may be associated with specific operational environments, security controls form part of broader governance frameworks, and business capabilities contribute to larger organisational objectives. By explicitly modelling containment and membership, the SKM preserves the organisational boundaries and contextual grouping of semantic entities, enabling more effective knowledge navigation and contextual reasoning. Structural relationships have also been designed to support semantic extensibility. As the CyberAId platform evolves, new semantic entities may be introduced to represent emerging technologies, regulatory developments, threat categories, or organisational concepts. The relationship model allows these new entities to be integrated into the existing semantic architecture through established hierarchical and compositional structures without requiring modifications to the overall conceptual organisation. This extensibility ensures that the SKM remains adaptable to future developments while preserving semantic integrity and interoperability. An important characteristic of structural relationships is that they are generally persistent. Unlike behavioural relationships, which often represent temporary operational interactions, structural relationships typically describe relatively stable organisational or conceptual associations. Although organisational structures, service architectures, or asset inventories may evolve over time, their structural relationships usually remain valid over extended periods and therefore provide a stable semantic foundation upon which more dynamic operational knowledge can be represented. This stability contributes to the long-term consistency of the SKM and facilitates incremental knowledge evolution without compromising the integrity of existing semantic structures.

Finally, structural relationships establish the organisational framework that supports all subsequent categories of semantic associations. Behavioural interactions, contextual dependencies, temporal evolution, and provenance information all operate upon the structural architecture defined by these relationships. Consequently, structural relationships should be regarded as the semantic backbone of the SKM, providing the stable organisational context required for advanced cybersecurity analysis, graph traversal, AI-assisted reasoning, and knowledge integration throughout the CyberAId ecosystem.

3.5.3 Behavioural Relationships

Behavioural relationships describe the dynamic interactions that occur between semantic entities throughout the cybersecurity lifecycle. While structural relationships establish the stable organisation of knowledge within the SKM, behavioural relationships capture how entities influence one another during operational activities, threat evolution, attack execution, vulnerability exploitation, incident response, and organisational defence. These relationships transform the SKM from a static representation of cybersecurity concepts into a dynamic semantic ecosystem capable of describing how cyber events unfold, evolve, and influence organisational operations over time. The Behavioural Relationship Model has been designed to represent operational interactions independently of any specific implementation technology or cybersecurity framework. Rather than encoding particular attack methodologies or incident response procedures, it provides a generic semantic abstraction through which operational behaviour can be represented consistently across heterogeneous cybersecurity scenarios. This technology-independent approach enables CyberAId to integrate threat intelligence, operational monitoring, defensive activities, and organisational response within a common semantic architecture while remaining adaptable to future analytical methodologies and implementation technologies. A central characteristic of behavioural relationships is the representation of causal interactions. Cybersecurity events rarely occur in isolation; instead, they form sequences of related activities in which one event enables, influences, or triggers another. A threat actor initiates an attack campaign, an attack employs specific techniques, those techniques exploit existing vulnerabilities, successful exploitation generates observable events, observations contribute to

detections, detections lead to investigations, and confirmed incidents activate defensive procedures. Behavioural relationships explicitly represent these causal dependencies, enabling CyberAId to preserve the logical progression of cybersecurity events throughout the knowledge ecosystem.

The Behavioural Relationship Model also supports the representation of operational workflows. Modern cybersecurity operations involve coordinated sequences of activities that extend beyond individual technical actions. Threat intelligence collection, vulnerability assessment, security monitoring, incident investigation, containment, eradication, recovery, post-incident analysis, and continuous improvement all represent interconnected operational processes rather than isolated events. Behavioural relationships capture these workflow-oriented interactions by describing how cybersecurity activities progress, interact, and exchange contextual knowledge during operational execution. This representation enables CyberAId to model complete cybersecurity workflows while preserving semantic continuity across different stages of cyber defence. Another important aspect concerns the modelling of offensive behaviour. Threat actors interact with infrastructures, attack techniques utilise offensive tools, malware communicates with command-and-control infrastructures, exploits compromise vulnerable assets, and attack campaigns evolve through multiple operational phases. Behavioural relationships describe these interactions by linking offensive entities according to their operational roles rather than their structural characteristics. This behavioural representation supports attack path reconstruction, adversarial behaviour analysis, campaign correlation, and threat intelligence enrichment while maintaining compatibility with recognised cybersecurity knowledge sources and adversarial modelling frameworks. Complementing offensive behaviour, the Behavioural Relationship Model also represents defensive behaviour. Security controls monitor organisational assets, detection mechanisms analyse operational observations, response procedures contain incidents, recovery activities restore affected services, and resilience measures strengthen organisational preparedness following operational disruptions. Behavioural relationships describe how these defensive capabilities interact throughout the cybersecurity lifecycle, enabling CyberAId to model coordinated cyber defence as a sequence of interconnected operational activities rather than independent technical mechanisms. This representation is particularly valuable for supporting orchestration, playbook execution, and AI-assisted response planning within future CyberAId components. Behavioural relationships further enable the semantic representation of knowledge evolution. Operational observations continuously enrich threat intelligence, newly discovered vulnerabilities modify exposure assessments, incident investigations generate additional evidence, defensive actions improve organisational preparedness, and lessons learned refine future detection capabilities. The SKM therefore captures not only operational activities but also the continuous exchange and refinement of knowledge resulting from those activities. Through these relationships, CyberAId supports adaptive cybersecurity in which operational experience progressively improves the quality, completeness, and contextual richness of the knowledge ecosystem. An equally important objective is the representation of interaction across semantic domains. Behavioural relationships are not confined to individual domains but frequently connect cybersecurity, financial, organisational, risk, and regulatory concepts. For example, an attack may disrupt a payment service, an operational incident may trigger regulatory reporting, a mitigation activity may reduce organisational risk, or a resilience exercise may improve future defensive capabilities. By explicitly modelling these interactions, the SKM enables CyberAId to reason across technical and business domains without fragmenting operational knowledge into separate conceptual models. The Behavioural Relationship Model also accommodates varying degrees of interaction complexity. Certain relationships describe direct interactions between two semantic entities, whereas others emerge from longer operational sequences involving multiple entities distributed across different semantic domains. Rather than limiting behavioural representation to pairwise associations, the proposed specification allows behavioural knowledge to be interpreted as interconnected operational chains capable of describing complete cybersecurity scenarios. This capability provides the conceptual basis for graph traversal, behavioural analytics, dependency analysis, attack simulation, and explainable AI reasoning across the CyberAId platform.

Finally, behavioural relationships establish the dynamic dimension of the SKM by describing how cybersecurity knowledge changes during operational execution. Together with the structural relationships

introduced in the previous subsection, they enable CyberAid to represent both the stable architecture of organisational knowledge and the dynamic interactions through which cybersecurity events continuously transform that knowledge. The following subsection extends this perspective further by introducing contextual relationships that associate operational cybersecurity behaviour with organisational, business, financial, and regulatory context.

3.5.4 Contextual Relationships

Contextual relationships provide the semantic mechanisms through which technical cybersecurity knowledge is associated with its broader organisational, financial, operational, risk, and regulatory context. While structural relationships organise semantic entities and behavioural relationships describe their operational interactions, contextual relationships establish the semantic associations that explain why those interactions are relevant within the organisational environment. They enable the SKM to enrich technical cybersecurity observations with business meaning, allowing CyberAid to interpret cybersecurity events according to organisational objectives, operational dependencies, financial services, resilience requirements, and regulatory obligations.

The primary objective of contextual relationships is to transform isolated cybersecurity information into context-aware organisational knowledge. Modern cybersecurity decision-making rarely depends solely on technical evidence; instead, it requires understanding how technical events influence business operations, organisational priorities, customer services, regulatory responsibilities, and enterprise resilience. Contextual relationships therefore associate cybersecurity entities with the operational environment in which they exist, enabling CyberAid to move beyond purely technical analysis and support multidimensional reasoning across the entire knowledge ecosystem. Unlike behavioural relationships, contextual relationships do not describe actions or event progression. Instead, they express semantic associations that characterise relevance, applicability, dependency, influence, business significance, organisational ownership, regulatory applicability, or strategic importance. These relationships remain meaningful even in the absence of active cybersecurity events because they describe the contextual environment within which operational interactions may eventually occur. Consequently, contextual relationships complement the dynamic nature of behavioural interactions by providing the stable business semantics required for informed decision-making. One of the most important categories of contextual relationships concerns the association between cybersecurity entities and business services. Threats, vulnerabilities, incidents, security controls, and defensive capabilities acquire significantly greater operational relevance when connected to the financial services they support or threaten. A vulnerability affecting a customer authentication platform, for example, is not merely associated with a technical component but also with identity verification services, digital banking operations, payment authorisation workflows, customer experience, and operational continuity. Contextual relationships preserve these associations explicitly, enabling CyberAid to evaluate cybersecurity events according to their organisational significance rather than their technical characteristics alone.

Contextual relationships also describe the association between cybersecurity knowledge and organisational assets and business processes. Technical observations frequently influence multiple organisational functions simultaneously, while individual business processes often depend upon numerous technological assets and external services. By explicitly representing these associations, the SKM captures the operational dependencies that exist between cybersecurity and organisational operations. This contextual representation enables impact analysis, business continuity assessment, and resilience-oriented reasoning without requiring separate business models disconnected from cybersecurity knowledge. Another important characteristic of contextual relationships is the representation of risk associations. Rather than directly calculating risk, these relationships establish the semantic context through which technical cybersecurity knowledge contributes to organisational risk scenarios. Threats may be associated with particular categories of operational risk, vulnerabilities may contribute to exposure assessments, incidents may influence business impact evaluations, and defensive capabilities may reduce residual risk. These contextual associations preserve the conceptual relationships necessary for downstream risk analysis while maintaining the separation between semantic representation and analytical assessment methodologies. The proposed specification further introduces contextual relationships linking

cybersecurity knowledge with governance and regulatory concepts. Security controls may support specific compliance obligations, incidents may require regulatory notification, organisational assets may fall within the scope of sector-specific legislation, and operational processes may be governed by multiple supervisory requirements. These relationships enable CyberAid to contextualise cybersecurity activities according to applicable governance frameworks without embedding regulatory logic directly into technical cybersecurity entities. This separation improves semantic clarity while facilitating explainable compliance reasoning and regulatory traceability. Contextual relationships also support the representation of criticality and prioritisation. Not every cybersecurity observation possesses equal organisational significance. Certain assets support essential business services, specific incidents affect critical infrastructures, and particular vulnerabilities expose strategically important operational capabilities. Contextual relationships therefore associate semantic entities with concepts describing business criticality, operational importance, service priority, resilience objectives, organisational value, and strategic relevance. These contextual associations enable CyberAid to support intelligent prioritisation of cybersecurity activities while preserving the organisational rationale behind prioritisation decisions. An equally important objective of the Contextual Relationship Model is the preservation of semantic interoperability across heterogeneous knowledge domains. Financial institutions manage information originating from cybersecurity platforms, governance systems, operational repositories, business process documentation, regulatory frameworks, and organisational knowledge bases. Contextual relationships provide the semantic bridges that enable these heterogeneous information sources to coexist within a unified SKM without requiring identical internal data structures or terminology. This interoperability enables CyberAid to integrate technical and business knowledge while maintaining conceptual consistency across the complete semantic ecosystem. The Contextual Relationship Model has also been designed to support AI-assisted contextual reasoning. Intelligent analytical systems require significantly more than isolated technical events to produce reliable recommendations or meaningful explanations. Contextual relationships enable AI components to associate technical observations with organisational priorities, business dependencies, resilience objectives, governance obligations, historical incidents, and operational context. Consequently, graph-based reasoning, Retrieval-Augmented Generation (RAG), intelligent cyber agents, and decision-support mechanisms can retrieve semantically enriched knowledge that reflects both technical reality and organisational relevance. This capability substantially improves explainability, contextual awareness, and decision quality across AI-enabled CyberAid components.

Finally, contextual relationships complete the semantic interpretation layer of the SKM by connecting operational cybersecurity knowledge with the wider organisational environment in which financial institutions operate. Together with structural and behavioural relationships, they establish a multidimensional semantic architecture capable of representing not only how cybersecurity events occur but also why they matter to the organisation, how they influence operational resilience, and how they contribute to strategic decision-making. The following subsection extends this semantic architecture by introducing temporal, provenance, and lifecycle relationships, enabling the SKM to represent the evolution, trustworthiness, and historical progression of cybersecurity knowledge over time.

3.5.5 Temporal and Lifecycle Relationships

Temporal and Lifecycle Relationships define the semantic associations that describe the evolution, validity, succession, and historical progression of knowledge within the CyberAid SKM. While structural relationships organise knowledge, behavioural relationships describe operational interactions, and contextual relationships provide organisational meaning, temporal relationships capture how cybersecurity knowledge changes over time. These relationships enable the SKM to preserve historical context, maintain traceability between successive knowledge states, support continuous intelligence updates, and represent the lifecycle of semantic entities throughout their existence.

Cybersecurity knowledge is inherently dynamic. Threat actors continuously modify their tactics, vulnerabilities progress from disclosure to remediation, attack campaigns evolve through multiple operational phases, indicators of compromise become obsolete, defensive controls mature over time, and regulatory frameworks undergo periodic revision. Consequently, the SKM cannot represent knowledge as a static collection of facts. Instead, it models cybersecurity information as an evolving semantic ecosystem

in which entities and their relationships continuously adapt to newly acquired intelligence, operational observations, organisational changes, and regulatory developments. A fundamental objective of temporal relationships is the preservation of chronological continuity. Semantic entities frequently pass through multiple states during their operational lifecycle, each representing a different stage of maturity, validity, or operational relevance. For example, a vulnerability may initially be reported, subsequently verified, later associated with exploit availability, eventually remediated, and finally archived as historical knowledge. Similarly, a cyber threat campaign may emerge, expand, adapt its behaviour, and ultimately become inactive. Temporal relationships explicitly connect these successive stages, enabling CyberAid to maintain a coherent historical narrative rather than isolated snapshots of cybersecurity information. The Temporal and Lifecycle Relationship Model also supports the representation of knowledge versioning. Semantic entities may undergo refinement as additional intelligence becomes available, analytical conclusions evolve, or organisational information changes. Rather than replacing previous knowledge, the SKM preserves relationships between successive versions, enabling analysts and automated reasoning mechanisms to understand how knowledge has evolved over time. This version-aware representation contributes to explainability, reproducibility, and historical analysis while supporting continuous refinement of cybersecurity intelligence.

Another important characteristic concerns the representation of knowledge validity. Cybersecurity information often possesses limited temporal relevance. Indicators of compromise may expire, vulnerabilities may cease to be exploitable following remediation, attack infrastructure may become inactive, certificates may be revoked, regulatory requirements may be superseded, and organisational assets may be decommissioned. Temporal relationships therefore associate semantic entities with validity periods, activation conditions, retirement events, review cycles, and expiration states. This representation enables CyberAid to distinguish between active operational knowledge and historical information while preventing obsolete intelligence from influencing future analytical processes. Temporal relationships further capture operational lifecycle transitions. Cybersecurity entities evolve through well-defined stages during their existence, and these transitions often carry important semantic meaning. Incidents progress from initial detection to investigation, containment, recovery, and closure. Defensive controls move from planning to implementation, validation, deployment, maintenance, and eventual replacement. Risk scenarios may be identified, assessed, mitigated, monitored, and periodically reassessed. Representing these lifecycle transitions explicitly enables CyberAid to model cybersecurity operations as continuously evolving processes rather than static organisational states. The proposed specification also accommodates the representation of event sequencing. Certain cybersecurity activities occur in strict chronological order and must therefore preserve temporal dependencies between individual observations or operational actions. Attack phases, incident response procedures, investigative workflows, resilience exercises, and regulatory reporting activities all involve ordered sequences whose temporal progression contributes to their semantic interpretation. Temporal relationships preserve these ordered dependencies without embedding procedural logic directly into the semantic entities themselves, thereby maintaining conceptual clarity while supporting chronological reasoning. An equally important objective of the Temporal and Lifecycle Relationship Model is the preservation of historical traceability. Operational decisions frequently require analysts to reconstruct previous cybersecurity situations, understand how knowledge evolved, verify historical evidence, or explain earlier analytical conclusions. By explicitly linking historical observations, intelligence updates, assessment revisions, remediation activities, and governance decisions, the SKM enables CyberAid to reconstruct complete knowledge histories while preserving semantic integrity across successive updates. This historical traceability significantly strengthens auditability, forensic analysis, and explainable AI capabilities throughout the platform. Temporal relationships additionally support continuous knowledge enrichment. Newly acquired intelligence rarely invalidates existing knowledge entirely; instead, it typically extends, refines, confirms, or supersedes previous observations. Threat intelligence reports may corroborate earlier findings, incident investigations may reveal additional attack stages, vulnerability disclosures may expand known exposure conditions, and regulatory updates may introduce revised compliance obligations. Temporal relationships preserve these refinement processes by explicitly associating successive knowledge states, thereby enabling the SKM to evolve incrementally while maintaining continuity between historical and current knowledge. Because CyberAid

aims to support intelligent reasoning over continuously evolving cybersecurity ecosystems, temporal relationships also contribute directly to AI-assisted analysis. Knowledge evolution, historical trends, lifecycle transitions, and temporal dependencies provide valuable contextual information for graph analytics, predictive reasoning, anomaly detection, trend analysis, Retrieval-Augmented Generation (RAG), and intelligent cyber agents. By preserving temporal semantics explicitly, the SKM enables AI components to reason not only about the current state of cybersecurity knowledge but also about its historical evolution and expected future development.

Finally, Temporal and Lifecycle Relationships complete the dynamic dimension of the SKM by introducing explicit representations of knowledge evolution, validity, historical continuity, and lifecycle progression. Together with the structural, behavioural, and contextual relationship models, they enable CyberAid to maintain a living knowledge ecosystem in which cybersecurity information evolves continuously while preserving semantic consistency, historical traceability, and analytical reliability across the complete lifecycle of organisational cybersecurity knowledge.

3.5.6 Cross-Domain Relationship Patterns

The preceding subsections have introduced the principal categories of semantic relationships that collectively define the structural, behavioural, contextual, and temporal characteristics of the CyberAid SKM. Individually, these relationship categories describe specific aspects of cybersecurity knowledge representation. However, real-world cybersecurity situations rarely exist within the boundaries of a single semantic domain or relationship type. Instead, they emerge through complex interaction patterns that simultaneously involve technical entities, organisational assets, business processes, threat intelligence, operational observations, defensive capabilities, risk scenarios, and regulatory obligations. The purpose of the Cross-Domain Relationship Patterns is therefore to define how these heterogeneous semantic relationships cooperate to form coherent, end-to-end knowledge structures across the entire SKM. A fundamental design principle of the CyberAid SKM is that knowledge is inherently cross-domain. Threat intelligence does not terminate at the identification of threat actors, vulnerabilities do not exist independently of the assets they affect, incidents cannot be interpreted without operational context, and regulatory obligations cannot be evaluated without considering the cybersecurity events and organisational processes to which they apply. Consequently, the SKM avoids isolated semantic silos by explicitly defining relationship patterns that traverse multiple semantic domains while preserving semantic consistency, traceability, and contextual integrity.

Cross-domain relationship patterns describe the logical pathways through which cybersecurity knowledge evolves from isolated technical observations into comprehensive organisational intelligence. Rather than representing disconnected relationships, the SKM models complete semantic chains in which successive relationships progressively enrich the available knowledge. For example, threat intelligence describing an adversarial campaign may be associated with specific attack techniques, which in turn exploit identified vulnerabilities affecting organisational assets that support critical financial services. These affected services contribute to defined business processes, influence operational risk scenarios, activate regulatory obligations, and ultimately require the implementation of appropriate security controls and organisational response activities. Each individual relationship contributes a specific semantic meaning, while their composition provides a holistic representation of the cybersecurity situation. One of the most representative cross-domain patterns concerns the progression from cyber threat intelligence to organisational impact. External threat intelligence identifies malicious actors, campaigns, or indicators that are semantically associated with attack techniques and exploit strategies. These attack techniques target exploitable vulnerabilities residing within technological assets that support financial services and business operations. Once contextualised within the organisational environment, these technical observations become associated with operational dependencies, business processes, resilience objectives, and organisational risk scenarios. Finally, governance responsibilities, compliance obligations, and defensive controls are connected to the same knowledge chain, enabling CyberAid to interpret a single cybersecurity event simultaneously from technical, operational, business, and regulatory perspectives.

A complementary relationship pattern captures the operational defence lifecycle. Security monitoring capabilities generate observations that contribute to detection activities, which may initiate investigations leading to incident confirmation. Confirmed incidents activate response procedures, containment measures, recovery activities, and post-incident analysis. The resulting knowledge is subsequently associated with threat intelligence repositories, defensive knowledge, organisational learning processes, and governance activities. This feedback mechanism illustrates how operational cybersecurity continuously enriches the SKM, enabling future analyses to benefit from accumulated organisational experience while preserving semantic continuity across successive operational cycles. Cross-domain relationship patterns also support risk-aware knowledge integration. Rather than calculating cyber risk directly, the SKM represents the semantic pathways through which technical events influence organisational exposure and resilience. Threat actors, vulnerabilities, operational incidents, business assets, critical services, defensive capabilities, governance activities, and regulatory requirements collectively contribute to the formation of contextual risk scenarios. These semantic pathways provide downstream CyberAid components with the structured knowledge required to perform cyber-risk assessment, resilience evaluation, decision support, and AI-assisted reasoning without embedding risk calculation methodologies within the semantic model itself.

Another important pattern concerns the integration of regulatory and governance knowledge with operational cybersecurity. Security controls are semantically associated with the regulatory obligations they support, organisational responsibilities are linked to operational assets and business processes, and incident handling activities may activate reporting obligations, audit requirements, or evidence preservation procedures. Rather than introducing a separate compliance model, the SKM integrates governance knowledge directly into the operational cybersecurity context through explicit semantic relationship chains. This integration enables explainable compliance assessment while preserving a unified representation of organisational cybersecurity knowledge. Cross-domain relationship patterns further facilitate knowledge interoperability between heterogeneous information sources. Cyber threat intelligence feeds, vulnerability repositories, organisational asset inventories, security monitoring platforms, governance documentation, regulatory frameworks, and business process repositories each contribute different perspectives on organisational cybersecurity. The SKM harmonises these heterogeneous sources by connecting them through common semantic relationship patterns rather than requiring identical internal data structures or information models. This capability enables CyberAid to integrate diverse knowledge sources into a single semantic ecosystem while maintaining conceptual consistency across all participating domains. An equally significant characteristic of the proposed relationship patterns is their suitability for graph-based reasoning and AI-assisted knowledge exploration. Because semantic entities are interconnected through explicit cross-domain relationship chains, graph analytics can traverse organisational dependencies, identify attack propagation paths, correlate operational incidents with business impact, analyse defensive effectiveness, and explore regulatory implications within a unified knowledge graph. Similarly, Retrieval-Augmented Generation (RAG), intelligent cyber agents, and explainable AI systems can retrieve semantically connected information rather than isolated technical artefacts, thereby producing more contextual, transparent, and operationally meaningful analytical outputs. The Cross-Domain Relationship Patterns have been intentionally defined at the conceptual level, ensuring compatibility with multiple implementation technologies and future platform evolution. Whether realised through ontology frameworks, property graphs, knowledge graphs, semantic databases, or hybrid AI-driven architectures, the underlying semantic patterns remain unchanged. This abstraction enables CyberAid to preserve the conceptual integrity of the SKM while allowing implementation decisions to evolve throughout subsequent development activities.

Ultimately, the Cross-Domain Relationship Patterns unify the entire semantic architecture of the SKM. By integrating structural, behavioural, contextual, and temporal relationships into coherent knowledge pathways, the model enables CyberAid to represent cybersecurity as an interconnected organisational ecosystem rather than a collection of isolated technical concepts. This integrated semantic architecture provides the conceptual foundation upon which CTI integration, knowledge representation, graph analytics,

explainable AI, cyber-risk reasoning, and operational decision support are subsequently realised throughout the CyberAid platform.

3.5.7 Relationship type catalogue

The following indicative catalogue defines the minimum typed relationships used by the SKM. Domain identifies the class or classes permitted as the source of a relationship, while range identifies the class or classes permitted as its target. Cardinality indicates how many source and target instances may participate in the relationship and may be refined by future implementation profiles without changing the relationship's semantic meaning.

Table 4 - Relationship Catalogue

Relationship	Domain	Range	Cardinality	Example
attributed_to	Campaign, Indicator, Malware	ThreatActor	many-to-many	Campaign attributed to ThreatActor
uses	ThreatActor, Campaign	AttackTechnique, Malware, Tool	many-to-many	Campaign uses AttackTechnique
exploits	AttackTechnique	Vulnerability	many-to-many	Technique exploits Vulnerability
affects	Vulnerability, Incident	Asset	many-to-many	Vulnerability affects Asset
observed_on	Observation, Indicator	Asset	many-to-many	Indicator observed on Asset
supports	Observation, SourceArtefact	Finding	many-to-many	Observation supports Finding
mitigates	SecurityControl	Vulnerability, AttackTechnique, Risk	many-to-many	Control mitigates Vulnerability
depends_on	Asset, BusinessProcess	Asset, Service	many-to-many	BusinessProcess depends on Service
subject_to	Asset, BusinessProcess, Organisation	RegulatoryRequirement	many-to-many	Service subject to regulatory requirement
derived_from	KnowledgeElement	SourceArtefact, KnowledgeElement	many-to-many	Indicator derived from CTI report

Relationship	Domain	Range	Cardinality	Example
same_as	KnowledgeElement	KnowledgeElement	many-to-many	SKM vulnerability corresponds to CVE record
superseded_by	KnowledgeElement	KnowledgeElement	zero-or-one successor	Requirement superseded by a newer version
owns	Organization	Asset	One-to-many	One Organisation owns multiple Assets, while each Asset has one accountable owner

The following simplified example illustrates how the principal SKM entities may be connected:

ThreatActor - uses→ AttackTechnique - exploits→ Vulnerability - affects→ Asset ←depends_on - BusinessProcess - subject_to→ RegulatoryRequirement

SecurityControl —mitigates→ Vulnerability

For example, a threat actor may use an attack technique that exploits a vulnerability affecting an asset on which a payment business process depends. A security control may mitigate the vulnerability, while the business process remains linked to the applicable regulatory requirement.

3.6 Metadata, Provenance and Governance Attributes

The semantic entities and relationships defined throughout the previous sections establish the conceptual structure of the CyberAId SKM. However, a semantic knowledge model intended to support cybersecurity operations, threat intelligence integration, graph analytics, and AI-assisted reasoning requires more than an accurate representation of concepts and their associations. It must also preserve information describing the origin, quality, reliability, validity, ownership, evolution, and governance of the knowledge it contains. Consequently, the SKM incorporates a complementary metadata and governance layer that enables cybersecurity knowledge to be interpreted not only according to its semantic meaning but also according to its trustworthiness, operational relevance, and lifecycle status.

Metadata provides the descriptive information necessary to characterise semantic entities and relationships beyond their conceptual definition. Provenance captures the origin and history of cybersecurity knowledge, enabling analysts and automated reasoning mechanisms to determine where information originated, how it has evolved, and under which conditions it should be trusted. Governance attributes establish the organisational policies through which knowledge is managed, protected, reviewed, shared, and maintained throughout its lifecycle. Collectively, these complementary mechanisms transform the SKM from a semantic representation into a trustworthy knowledge ecosystem capable of supporting explainable cybersecurity analysis and evidence-driven decision-making. The metadata model has been designed according to the principles of transparency, traceability, interoperability, consistency, technology independence, and extensibility. Metadata attributes are therefore defined conceptually rather than implementation-specifically, allowing the SKM to preserve identical semantic meaning regardless of the

technologies used for storage, graph representation, ontology management, or AI-enabled knowledge processing. This abstraction also enables future CyberAId components to extend metadata definitions without affecting the underlying semantic architecture.

The following subsections define the principal metadata categories adopted by the SKM, including descriptive metadata, provenance information, knowledge lifecycle management, governance mechanisms, access control principles, explainability, and auditability. Together, these elements establish the knowledge governance layer required to ensure that semantic information remains reliable, traceable, maintainable, and suitable for continuous operational use throughout the CyberAId platform.

3.6.1 Metadata Model

Metadata constitute the descriptive layer of the CyberAId SKM, providing the contextual information required to interpret, manage, validate, and govern semantic knowledge throughout its operational lifecycle. While semantic entities and relationships define the conceptual structure of cybersecurity knowledge, metadata describe the characteristics of that knowledge, enabling both human analysts and automated reasoning mechanisms to evaluate its origin, quality, relevance, validity, and operational suitability. Consequently, the Metadata Model establishes the descriptive foundation upon which the trustworthiness and long-term maintainability of the SKM are built. Within the proposed specification, metadata are treated as an integral component of every semantic entity and relationship rather than as external annotations. Every concept represented within the SKM, including threat actors, attack techniques, vulnerabilities, organisational assets, business processes, incidents, security controls, risk scenarios, regulatory obligations, and the semantic relationships connecting them, may be associated with descriptive metadata that characterise its operational context and management requirements. This approach ensures that semantic knowledge is accompanied by sufficient contextual information to support informed interpretation, consistent maintenance, and transparent reasoning across the CyberAId ecosystem. The Metadata Model has been designed according to the principles of consistency, technology independence, interoperability, extensibility, and minimal semantic redundancy. Metadata attributes are defined conceptually rather than according to the implementation requirements of specific graph databases, ontology languages, or cybersecurity platforms. As a result, the same metadata concepts may be implemented through ontology annotations, graph properties, structured documents, relational repositories, or hybrid knowledge management technologies while preserving identical semantic meaning. This abstraction ensures that the metadata model remains stable even as implementation technologies evolve throughout the lifecycle of the CyberAId platform.

Table 5 - Metadata and Provenance

Field	Datatype	Requirement	Purpose
element_id	identifier	Required	Stable canonical identity
element_type	controlled term	Required	SKM class
created_at	datetime	Required	Initial creation time
modified_at	datetime	Required	Latest semantic update
version	string	Required	Element version

Field	Datatype	Requirement	Purpose
source_reference	reference or list	Required	Authoritative source locator
provenance_reference	reference or list	Required	Acquisition and transformation record
external_identifier	identifier or list	Conditional	Identifier assigned by an external authority
confidence	integer 0–100	Optional	Confidence in correctness
classification	controlled term	Required	CyberAId sensitivity level
tlp_marking	controlled term	Conditional	Original TLP dissemination marking
handling_constraints	list	Conditional	Source, legal, contractual or privacy restrictions
knowledge_partition	identifier or list	Required for retrievable elements	Logical knowledge partition
semantic_domain	identifier or list	Required	Owning semantic domain
graph_scope	identifier or list	Conditional	Permitted semantic traversal scope
valid_from	datetime	Optional	Beginning of validity
valid_until	datetime	Optional	End of validity
lifecycle_state	controlled term	Required	Current knowledge state
owner	reference	Required	Accountable owner
steward	reference	Optional	Operational steward

Field	Datatype	Requirement	Purpose
revoked	boolean	Required, default false	Revocation indicator

A primary objective of the Metadata Model is the provision of descriptive metadata that facilitate the identification and interpretation of semantic knowledge. Every semantic entity may therefore be associated with human-readable identifiers, names, descriptions, aliases, classifications, semantic categories, domain assignments, and external references. These descriptive attributes improve knowledge discoverability, facilitate interoperability with external cybersecurity repositories, and support efficient navigation of the SKM without altering the underlying semantic structure. The Metadata Model also supports the representation of operational metadata, describing the current status and practical characteristics of semantic knowledge. Such metadata may indicate whether a knowledge element is active, deprecated, archived, under review, or pending validation, while also identifying its operational scope, intended usage, associated organisational context, and applicable semantic domains. This information enables CyberAid components to distinguish operational knowledge suitable for real-time reasoning from historical or experimental information maintained primarily for reference or analytical purposes. Another important category comprises classification metadata, which organise semantic entities according to predefined conceptual dimensions without modifying their intrinsic semantic relationships. Classification attributes may describe the type of cybersecurity knowledge represented, the associated semantic domain, the business context, the relevant operational function, the applicable regulatory area, or the intended analytical purpose. These metadata enable flexible filtering, grouping, and retrieval of knowledge while supporting domain-specific views of the SKM tailored to different analytical or operational requirements.

The proposed specification further accommodates reference metadata that establish semantic connections with external knowledge sources. Cybersecurity entities frequently originate from recognised repositories, threat intelligence platforms, vulnerability databases, regulatory documents, or standardisation initiatives. Rather than duplicating this external information, the Metadata Model enables semantic entities to reference authoritative external identifiers, documentation, standards, taxonomies, or intelligence reports. These references improve interoperability, facilitate future synchronisation with external knowledge sources, and preserve semantic consistency across heterogeneous cybersecurity ecosystems. The Metadata Model has also been designed to support AI-assisted knowledge utilisation. Intelligent analytical components require rich contextual information to retrieve relevant knowledge, explain analytical results, and distinguish authoritative information from lower-confidence observations. Metadata describing semantic categories, operational context, domain membership, descriptive attributes, and external references significantly improve graph traversal, semantic search, Retrieval-Augmented Generation (RAG), knowledge retrieval, and explainable AI by providing additional contextual cues beyond the semantic relationships themselves. Consequently, metadata become an essential component of knowledge discovery and contextual reasoning throughout the CyberAid platform. An equally important objective of the Metadata Model is to ensure semantic interoperability across heterogeneous information environments. Because CyberAid integrates knowledge originating from cyber threat intelligence platforms, vulnerability repositories, monitoring infrastructures, financial systems, governance documentation, and regulatory frameworks, metadata provide the common descriptive vocabulary required to harmonise these diverse information sources within a unified SKM. This interoperability enables consistent interpretation of semantic knowledge regardless of its origin while preserving the flexibility necessary to accommodate future information sources and evolving cybersecurity standards. The Metadata Model intentionally distinguishes descriptive metadata from the provenance, governance, and lifecycle information introduced in the following subsections. While metadata provide the descriptive characteristics of semantic knowledge, provenance explains its origin, lifecycle attributes describe its evolution over time, and governance attributes define the organisational mechanisms through which knowledge is managed and controlled.

Maintaining this separation improves conceptual clarity while enabling each category of knowledge descriptors to evolve independently without affecting the underlying semantic architecture.

Ultimately, the Metadata Model establishes the descriptive foundation of the CyberAid SKM by ensuring that every semantic entity and relationship can be consistently identified, interpreted, categorised, and reused throughout the platform. By enriching semantic knowledge with structured descriptive information, the model enhances interoperability, semantic discovery, AI-assisted reasoning, graph-based analytics, and long-term knowledge maintainability while preparing the transition towards the provenance and trust mechanisms described in the following subsection.

3.6.1.1 Identifier and Addressability Requirements

The canonical SKM identifier shall follow the indicative form “urn:cyberaid:skm:<element-type>:<identifier>”. Where an authoritative external identifier exists, including a STIX, CVE, CWE, CAPEC or MITRE ATT&CK identifier, it shall be preserved in an “external_identifier” attribute and used for reconciliation where appropriate. Internally derived identifiers should be generated deterministically from the authoritative source identifier and source namespace where this is feasible; otherwise, a UUID-based identifier shall be used. Deduplication and entity reconciliation shall not delete historical identifiers. Deprecated or merged identifiers shall remain resolvable through “same_as”, “superseded_by” or “merged_into” relationships. Revocation, archival and retirement shall modify the lifecycle state of an element without reassigning its canonical identifier. A retrieval result shall expose both the canonical SKM identifier and the source locator required to resolve its supporting evidence.

3.6.2 Provenance and Trust

The effectiveness of the CyberAid SKM depends not only on the semantic representation of cybersecurity knowledge but also on the ability to determine where that knowledge originates, how it has evolved, and to what extent it can be considered reliable. Cyber Threat Intelligence is inherently heterogeneous, originating from multiple internal and external sources that differ significantly in quality, authority, timeliness, completeness, and operational relevance. Consequently, the SKM incorporates an explicit Provenance and Trust Model that enables semantic knowledge to be evaluated according to its origin, supporting evidence, confidence, and overall trustworthiness before being utilised for analysis, reasoning, or decision support. Within the proposed specification, provenance describes the complete history associated with a semantic knowledge element, including its source of origin, acquisition pathway, transformation history, validation activities, and subsequent enrichment processes. Rather than treating cybersecurity knowledge as isolated information, the SKM preserves explicit provenance information throughout the entire lifecycle of every semantic entity and relationship. This approach enables CyberAid to reconstruct the complete knowledge lineage of a semantic object, thereby supporting transparency, reproducibility, explainability, and long-term knowledge governance.

A fundamental principle of the Provenance Model is that semantic knowledge should remain traceable to its original evidence whenever possible. Threat intelligence observations, vulnerability information, security events, regulatory knowledge, organisational assets, and semantic mappings may all originate from different authoritative sources. Instead of replacing or obscuring these origins during semantic integration, the SKM maintains explicit references to the original information providers, allowing analysts and automated reasoning mechanisms to determine the evidential basis supporting every semantic assertion. This provenance-aware representation significantly improves analytical transparency while reducing the risk of unsupported or unverifiable knowledge propagation.

Closely associated with provenance is the concept of trust, which represents the degree of confidence that may reasonably be assigned to semantic knowledge based on available evidence, source credibility, validation status, and contextual consistency. Rather than assuming that all information possesses equal reliability, the SKM recognises that cybersecurity knowledge frequently exhibits varying levels of certainty. Intelligence reports may contain preliminary observations, vulnerability disclosures may initially lack technical validation, threat actor attribution may remain uncertain, and organisational knowledge may evolve as additional operational evidence becomes available. The Trust Model therefore enables CyberAid to preserve these uncertainties explicitly instead of presenting all knowledge as equally authoritative. The

proposed specification distinguishes between several complementary dimensions of trust. Source trustworthiness reflects the credibility of the organisation or system providing the information, while information confidence expresses the degree of certainty associated with the specific knowledge element itself. Additional trust dimensions may capture validation status, corroboration by independent sources, analytical maturity, and consistency with existing semantic knowledge. Maintaining these dimensions independently allows CyberAId to differentiate between highly trusted information obtained from authoritative sources and preliminary observations requiring additional validation before influencing operational decision-making. The Provenance and Trust Model also accommodates the representation of knowledge transformation history. Cybersecurity information rarely enters the SKM in its original form. Threat intelligence may be normalised, vulnerabilities may be enriched with contextual information, semantic mappings may associate external entities with internal concepts, and organisational knowledge may be refined through successive analytical processes. Rather than obscuring these transformations, the SKM preserves semantic relationships describing how knowledge has been processed, enriched, correlated, or interpreted throughout its lifecycle. This transformation history enables analysts to distinguish between original observations and derived knowledge while improving explainability and supporting verification activities. Another important characteristic concerns the preservation of evidence associations. Semantic knowledge should not exist independently of the evidence supporting it. Threat intelligence reports, vulnerability disclosures, monitoring observations, incident investigations, forensic artefacts, regulatory documentation, organisational records, and analytical assessments collectively provide the evidence upon which semantic assertions are based. The SKM therefore explicitly associates semantic entities with supporting evidence wherever available, enabling CyberAId to justify analytical conclusions through verifiable information rather than implicit assumptions. This evidence-oriented approach significantly strengthens explainable AI, operational transparency, and audit readiness. The Provenance and Trust Model further supports multi-source knowledge correlation. Identical or related cybersecurity knowledge may originate from several independent intelligence providers, monitoring systems, organisational repositories, or regulatory documents. Rather than treating these observations as duplicate information, the SKM preserves their independent provenance while enabling semantic correlation across multiple evidence sources. This capability allows CyberAId to identify corroborating intelligence, resolve inconsistencies, strengthen confidence assessments, and maintain a richer contextual understanding of cybersecurity situations without sacrificing traceability. An equally important objective is the preservation of knowledge integrity. Provenance information enables every semantic entity to be associated with identifiable acquisition, validation, and modification activities, thereby reducing the likelihood of unauthorised or undocumented alterations. Combined with versioning and governance mechanisms introduced in subsequent sections, provenance relationships contribute to the long-term integrity of the SKM by ensuring that every significant knowledge modification remains attributable, reviewable, and historically traceable. Because CyberAId aims to support AI-assisted cybersecurity analysis, provenance and trust information also provide essential context for explainable reasoning. AI systems should not only produce analytical conclusions but also explain the evidence and reasoning upon which those conclusions are based. Provenance metadata enable intelligent components to identify the origin of retrieved knowledge, assess the confidence associated with competing intelligence sources, explain the analytical pathways followed during reasoning, and communicate uncertainty where appropriate. This capability significantly improves transparency, analyst confidence, and organisational acceptance of AI-assisted cybersecurity recommendations.

Table 6 - Confidence Scale

Label	Numeric range	Canonical display value
UNSPECIFIED	Field absent	Not applicable

Label	Numeric range	Canonical display value
NONE	0	0
LOW	1-29	15
MEDIUM	30-69	50
HIGH	70-100	85

The SKM confidence field follows the STIX 2.1 confidence property and uses an integer from 0 to 100. Absence of the field means that confidence is unspecified and shall not be interpreted as zero.

Thus, the Provenance and Trust Model establishes the trust layer of the CyberAId SKM. By explicitly representing knowledge origin, evidence, confidence, transformation history, and source credibility, the model ensures that cybersecurity knowledge remains transparent, verifiable, explainable, and suitable for operational decision-making. This trust-oriented semantic architecture strengthens graph analytics, CTI integration, explainable AI, cyber-risk reasoning, and regulatory accountability while preparing the transition towards the lifecycle management mechanisms introduced in the following subsection.

3.6.3 Knowledge Lifecycle

The CyberAId SKM has been designed to support continuously evolving cybersecurity environments in which knowledge is generated, refined, validated, shared, and retired throughout the operational lifecycle of the platform. Cyber Threat Intelligence, organisational knowledge, regulatory information, and cybersecurity observations are inherently dynamic, requiring continuous adaptation as new threats emerge, infrastructures evolve, business processes change, and regulatory obligations are updated. Consequently, the SKM incorporates an explicit Knowledge Lifecycle Model that governs the management and evolution of semantic knowledge while preserving consistency, traceability, and operational reliability across the complete knowledge ecosystem.

The Knowledge Lifecycle defines the organisational processes through which semantic knowledge is created, maintained, validated, updated, reused, archived, and ultimately retired. Rather than treating cybersecurity knowledge as static information, the proposed specification recognises that every semantic entity and relationship progresses through multiple lifecycle stages during its existence. These stages describe the operational maturity and management status of knowledge independently of its semantic meaning, enabling CyberAId to maintain an accurate and trustworthy knowledge repository despite the continuously changing cybersecurity landscape. The lifecycle begins with the knowledge acquisition and creation stage. New semantic knowledge may originate from external cyber threat intelligence feeds, vulnerability repositories, regulatory publications, organisational asset inventories, operational monitoring infrastructures, incident investigations, security assessments, or expert analysis. Regardless of its origin, newly acquired knowledge is incorporated into the SKM only after being represented according to the semantic structures defined throughout the previous sections. This initial stage establishes the conceptual identity of new knowledge while preserving its provenance and associated descriptive metadata.

Following acquisition, semantic knowledge enters a validation and enrichment phase. Newly created entities are reviewed for semantic consistency, completeness, interoperability, and contextual relevance before becoming part of the operational knowledge ecosystem. During this stage, additional contextual information may be incorporated, relationships with existing entities may be established, provenance

information may be completed, and external references may be associated with the new knowledge. Validation activities ensure that newly acquired information satisfies the quality expectations of the SKM while maintaining conceptual consistency across heterogeneous knowledge sources. Once validated, knowledge becomes available for operational utilisation. During this stage, semantic entities and relationships support graph analytics, cyber threat intelligence correlation, semantic search, Retrieval-Augmented Generation (RAG), AI-assisted reasoning, operational decision support, cyber-risk analysis, compliance assessment, and organisational resilience evaluation. Operational use does not represent the final stage of the lifecycle but rather the period during which semantic knowledge actively contributes to the analytical capabilities of the CyberAid platform while remaining subject to continuous monitoring and refinement. The Knowledge Lifecycle also accommodates continuous maintenance and refinement. Cybersecurity knowledge rarely remains unchanged after its initial incorporation into the SKM. Threat intelligence reports may be updated, vulnerability information may evolve, organisational assets may change ownership or operational status, regulations may be revised, and incident investigations may reveal additional contextual information. Rather than replacing existing semantic knowledge, the lifecycle model supports controlled refinement through incremental updates that preserve provenance, historical traceability, and semantic consistency. This approach ensures that knowledge continuously evolves without compromising the integrity of previously established semantic relationships.

An equally important lifecycle activity concerns knowledge review and quality assurance. Operational cybersecurity environments require periodic reassessment of stored knowledge to ensure its continued relevance, accuracy, and consistency. Review activities may identify obsolete intelligence, duplicated information, incomplete semantic mappings, inconsistent metadata, outdated regulatory references, or relationships requiring refinement. By incorporating systematic review mechanisms into the lifecycle, the SKM maintains long-term quality while reducing semantic drift and preserving the operational usefulness of cybersecurity knowledge. The lifecycle model further defines procedures for knowledge deprecation and retirement. Certain semantic entities gradually lose operational relevance as infrastructures are decommissioned, vulnerabilities are permanently remediated, threat campaigns cease activity, regulatory provisions are superseded, or business processes are redesigned. Rather than permanently deleting obsolete knowledge, the SKM distinguishes between active, deprecated, archived, and retired knowledge states. This distinction enables CyberAid to preserve valuable historical information for forensic analysis, trend identification, compliance evidence, and explainable AI while preventing obsolete knowledge from influencing current operational reasoning. Another significant aspect of the lifecycle concerns knowledge reuse and continuous learning. Lessons learned from incident investigations, resilience exercises, vulnerability assessments, cyber threat intelligence analysis, and operational experience should continuously improve the quality of future cybersecurity knowledge. The lifecycle therefore supports the systematic incorporation of newly generated organisational knowledge into the SKM, enabling CyberAid to evolve through operational experience rather than relying exclusively on external intelligence sources. This capability contributes directly to organisational learning, continuous improvement, and adaptive cyber resilience. The Knowledge Lifecycle has also been designed to support controlled knowledge evolution. Rather than permitting unrestricted modifications, every significant update is expected to preserve semantic consistency, maintain provenance information, and remain compatible with existing relationship structures and metadata definitions. This governance-oriented approach ensures that the SKM evolves in a controlled and predictable manner, reducing the likelihood of inconsistent semantic representations while supporting long-term maintainability across successive project phases. Because CyberAid aims to support trustworthy AI-assisted cybersecurity analysis, lifecycle information also contributes directly to knowledge transparency. Intelligent reasoning mechanisms should be capable of distinguishing newly acquired information from mature operational knowledge, recognising superseded knowledge elements, identifying recently validated intelligence, and understanding the historical evolution of semantic entities. Lifecycle attributes therefore provide valuable contextual information that improves explainability, supports analyst confidence, and enables more informed utilisation of cybersecurity knowledge during automated reasoning processes.

Eventually, the Knowledge Lifecycle establishes the operational management framework of the CyberAid SKM. By defining how semantic knowledge is created, validated, utilised, maintained, refined, archived, and retired, the lifecycle ensures that the SKM remains accurate, trustworthy, and operationally relevant despite the continuously changing nature of cybersecurity knowledge. Together with metadata and provenance, the lifecycle model completes the knowledge management foundations of the SKM while preparing the transition towards governance and access control mechanisms described in the following subsection.

3.6.4 Governance and Access Control

The long-term effectiveness of the CyberAid SKM depends not only on the quality and evolution of the knowledge it contains but also on the governance mechanisms through which that knowledge is managed, protected, and made available to authorised stakeholders. As the SKM integrates cybersecurity intelligence, organisational information, business knowledge, risk representations, and regulatory concepts originating from multiple internal and external sources, appropriate governance principles are required to ensure that semantic knowledge remains consistent, trustworthy, secure, and aligned with organisational policies throughout its operational lifecycle. Consequently, the SKM incorporates a Governance and Access Control Model that defines the organisational responsibilities, management principles, and access policies governing the creation, maintenance, sharing, and utilisation of semantic knowledge.

Knowledge governance within the SKM extends beyond traditional information management practices. Rather than focusing solely on data ownership or technical access permissions, the proposed model establishes a comprehensive governance framework covering semantic consistency, organisational accountability, knowledge stewardship, quality management, collaboration, and controlled information sharing. This governance-oriented perspective ensures that the evolution of semantic knowledge remains coordinated across the CyberAid ecosystem while supporting interoperability between the diverse technical and organisational components participating in the platform. A fundamental principle of the Governance Model is the establishment of clear ownership and stewardship responsibilities. Every semantic entity, relationship, metadata element, and associated knowledge artefact should be linked to identifiable organisational responsibilities describing ownership, maintenance, review, approval, and operational accountability. Knowledge ownership defines the organisational authority responsible for the correctness and long-term maintenance of semantic knowledge, while stewardship represents the operational responsibility for ensuring that knowledge remains complete, consistent, and aligned with the conceptual principles of the SKM. This distinction promotes accountability while facilitating coordinated collaboration between domain experts, cybersecurity analysts, developers, and governance stakeholders. The Governance Model also defines principles for knowledge contribution and modification. As cybersecurity knowledge continuously evolves through new threat intelligence, operational observations, regulatory updates, and organisational changes, controlled mechanisms are required to govern how semantic knowledge is introduced or modified. Rather than allowing unrestricted modifications, the SKM assumes that knowledge contributions follow defined governance procedures including review, validation, approval, and controlled publication. These procedures reduce the likelihood of inconsistent semantic representations, duplicate knowledge, or unauthorised modifications while preserving the conceptual integrity of the knowledge ecosystem.

An equally important aspect concerns the governance of knowledge quality. Maintaining high-quality semantic knowledge requires continuous monitoring of completeness, consistency, relevance, accuracy, and interoperability across all semantic domains. Governance activities therefore include periodic quality assessments, semantic consistency reviews, duplicate identification, metadata verification, relationship validation, and conceptual alignment with the evolving architecture of the SKM. These governance mechanisms complement the lifecycle processes introduced in the previous subsection by ensuring that operational evolution does not compromise semantic quality or conceptual coherence. The Governance Model further establishes principles for controlled knowledge sharing across organisational boundaries. CyberAid integrates knowledge originating from multiple stakeholders, cybersecurity platforms, organisational repositories, and external intelligence providers, each operating under different operational constraints, confidentiality requirements, and sharing policies. Consequently, semantic knowledge should only be disseminated according to predefined governance rules that respect ownership, intellectual

property, confidentiality obligations, licensing conditions, and applicable information-sharing agreements. This controlled sharing model promotes collaboration while preserving the integrity and confidentiality of sensitive cybersecurity knowledge. Access control within the SKM is defined as a governance mechanism regulating the visibility and utilisation of semantic knowledge rather than as a technical implementation of authentication or authorisation services. The model therefore specifies conceptual access policies describing which categories of knowledge may be viewed, modified, enriched, validated, or shared by different organisational roles according to operational responsibilities and governance requirements. Such policies remain independent of specific identity management technologies or access control implementations, allowing future CyberAid components to realise these principles through appropriate technical mechanisms without affecting the conceptual architecture of the SKM. The proposed specification recognises that semantic knowledge may exhibit different levels of sensitivity. Certain knowledge elements may represent publicly available threat intelligence, whereas others may contain sensitive organisational information, confidential incident records, proprietary analytical results, or restricted regulatory evidence. Governance attributes therefore support semantic classification according to information sensitivity, dissemination restrictions, handling requirements, and sharing permissions. This classification enables CyberAid to preserve appropriate confidentiality while maintaining semantic interoperability across information originating from multiple operational environments. The Governance and Access Control Model further recognises that semantic knowledge may be logically partitioned according to organisational, domain-specific, sensitivity, tenancy, provenance, or policy-related boundaries. Such partitioning enables different knowledge consumers to operate over authorised subsets of the overall CyberAid knowledge environment without implying unrestricted access to the complete SKM. The specific knowledge partitions available to a requesting component are therefore determined according to the authenticated security context and the applicable governance and authorisation policies at retrieval time.

TLP dissemination markings and organisational sensitivity classifications represent related but distinct governance dimensions [12]. A TLP marking defines the permitted dissemination boundary of the information and shall be retained as a residual handling constraint. The CyberAid classification attribute expresses the sensitivity level used by authorised platform consumers. Where both dimensions are present, the most restrictive applicable combination shall govern the use of the knowledge. A mapping shall never remove the original TLP restriction or override stricter legal, contractual, privacy, source-specific or organisational constraints.

Table 7 - Classification

CyberAid level	Rank	Meaning
PUBLIC	0	Information approved for unrestricted disclosure
INTERNAL	1	Information limited to the deploying organisation or an explicitly defined community
CONFIDENTIAL	2	Need-to-know information limited to authorised roles, clients or partners
RESTRICTED	3	Information limited to explicitly authorised recipients or processing contexts

Table 8 - TLP Mapping

Incoming marking	Default CyberAId level	Residual constraint
TLP:CLEAR	PUBLIC	Public disclosure permitted, subject to copyright and applicable law
TLP:GREEN	INTERNAL	Sharing limited to the defined community
TLP:AMBER	CONFIDENTIAL	Need-to-know sharing within the organisation and its clients
TLP:AMBER+STRICT	CONFIDENTIAL	Sharing limited to the recipient organisation
TLP:RED	RESTRICTED	Sharing limited to the individual recipients
TLP:WHITE	PUBLIC	Deprecated TLP 1.0 alias mapped to TLP:CLEAR while retaining the source marking
Unmarked	Organisation-defined default	No unrestricted dissemination until the applicable policy has been determined

These mappings define conservative semantic-normalisation defaults. They do not constitute an access-control decision. Runtime authorisation remains subject to the applicable Security Context and enforcement mechanism.

Another important governance objective concerns the management of semantic consistency across distributed knowledge contributors. As multiple organisations and domain experts may contribute to the evolution of the SKM, governance mechanisms are required to ensure that independently developed knowledge remains compatible with the common conceptual architecture. Shared modelling principles, controlled terminology, semantic review procedures, relationship validation, and metadata governance collectively contribute to preserving interoperability across distributed contributors while allowing controlled extension of the knowledge ecosystem as CyberAId evolves. The Governance Model also supports policy-driven knowledge management. Organisational policies governing cybersecurity, operational resilience, information sharing, knowledge retention, review frequency, and quality assurance may all influence how semantic knowledge is managed throughout its lifecycle. Rather than embedding organisational policies directly into semantic entities, the SKM associates governance policies with the relevant knowledge artefacts, enabling CyberAId to adapt governance practices according to evolving organisational or regulatory requirements without altering the semantic structure of the model itself.

Because the SKM is expected to support AI-assisted reasoning, governance mechanisms also contribute directly to trustworthy AI operation. Intelligent analytical components should operate only upon authorised, validated, and policy-compliant semantic knowledge while respecting organisational governance principles and information-sharing constraints. Governance attributes therefore provide essential contextual information enabling AI systems to understand which knowledge may be utilised for specific analytical purposes, under which conditions it may be shared, and how organisational governance

influences automated reasoning processes. This governance-aware approach strengthens transparency, operational trust, and responsible AI utilisation throughout the CyberAid platform.

Thus, the Governance and Access Control Model establishes the organisational management framework that enables the SKM to evolve in a controlled, secure, and accountable manner. By defining ownership, stewardship, contribution procedures, quality governance, controlled sharing, conceptual access policies, and policy-driven knowledge management, the model ensures that semantic knowledge remains consistent, trustworthy, interoperable, and aligned with organisational objectives throughout the operational lifecycle of CyberAid.

3.6.5 Explainability and Auditability

The CyberAid SKM has been designed to support cybersecurity analysis, threat intelligence integration, graph analytics, and AI-assisted decision support in a transparent, accountable, and evidence-driven manner. As cybersecurity operations increasingly rely on automated reasoning, semantic knowledge integration, and intelligent analytical services, the ability to understand, justify, and verify the conclusions produced by these mechanisms becomes essential. Consequently, the SKM incorporates Explainability and Auditability as fundamental governance principles, ensuring that semantic knowledge remains interpretable, traceable, verifiable, and suitable for both operational decision-making and regulatory oversight.

Within the context of the SKM, explainability refers to the capability of understanding how semantic knowledge contributes to analytical outcomes. Rather than presenting isolated conclusions, the model enables analysts, decision-makers, and AI-assisted components to identify the semantic entities, relationships, contextual associations, metadata, provenance information, and supporting evidence that collectively justify a particular interpretation or recommendation. This capability allows cybersecurity reasoning to remain transparent and understandable throughout the CyberAid platform, reducing dependence on opaque analytical processes and improving confidence in AI-assisted decision support. The explainability principles adopted by the SKM are founded upon the explicit semantic architecture established throughout this deliverable. Because cybersecurity knowledge is represented through well-defined entities, typed relationships, contextual associations, provenance information, and governance attributes, every analytical outcome can be associated with the semantic pathway that produced it. Rather than treating reasoning processes as isolated computational activities, CyberAid enables the reconstruction of semantic reasoning chains that describe how individual knowledge elements contributed to the final analytical conclusion. These semantic pathways significantly improve transparency while supporting expert validation of automated reasoning processes.

A central objective of explainability is the preservation of evidence-based reasoning. Analytical conclusions should not rely solely on inferred knowledge but should remain explicitly associated with the observations, intelligence reports, vulnerability disclosures, monitoring information, regulatory documentation, organisational records, and other evidence supporting the reasoning process. By maintaining these evidence associations, the SKM enables CyberAid to justify analytical outputs through verifiable semantic knowledge rather than implicit assumptions or opaque computational models. This evidence-oriented approach strengthens analyst confidence while facilitating independent verification of cybersecurity assessments. The Explainability Model also supports context-aware interpretation of cybersecurity knowledge. Technical cybersecurity observations frequently acquire different meanings depending on the organisational environment in which they occur. A vulnerability affecting a critical payment infrastructure may require substantially different interpretation than the same vulnerability affecting a non-critical internal service. By preserving contextual relationships linking technical observations with organisational assets, business processes, resilience objectives, and regulatory obligations, the SKM enables CyberAid to explain analytical conclusions according to both technical evidence and organisational significance. This multidimensional interpretation improves the operational usefulness of cybersecurity analysis while reducing ambiguity during decision-making. Closely related to explainability is the concept of auditability, which describes the capability to reconstruct, review, and verify the complete history of semantic knowledge and the analytical processes that utilise it. Auditability ensures that every significant semantic

entity, relationship, metadata element, provenance record, lifecycle event, governance action, and knowledge modification can be traced throughout its operational existence. This comprehensive traceability enables organisations to review historical decisions, reproduce analytical outcomes, validate semantic consistency, and demonstrate compliance with internal governance policies and applicable regulatory requirements. The Auditability Model further supports the preservation of complete knowledge lineage. Throughout the lifecycle of semantic knowledge, entities may be enriched, updated, validated, correlated, or superseded by new information. Rather than obscuring these modifications, the SKM preserves explicit records describing how knowledge has evolved, which governance activities have influenced its lifecycle, and which evidence has supported each significant modification. This historical continuity enables CyberAid to reconstruct previous knowledge states, investigate historical cybersecurity events, support forensic analysis, and maintain complete accountability for semantic knowledge evolution.

Another important aspect concerns the recording of governance and decision activities. The SKM not only preserves cybersecurity observations but also records the governance processes associated with semantic knowledge management. Validation decisions, review outcomes, approvals, quality assessments, policy updates, contribution activities, and lifecycle transitions collectively contribute to an auditable history describing how knowledge has been managed throughout its operational lifecycle. This governance-oriented audit trail strengthens organisational accountability while supporting transparent collaboration among distributed contributors and domain experts. The Explainability and Auditability Model has also been designed to support AI-assisted reasoning transparency. Intelligent analytical components should be capable of communicating not only their conclusions but also the semantic evidence, contextual relationships, provenance information, confidence assessments, and governance attributes supporting those conclusions. By preserving these complementary knowledge dimensions, the SKM enables Retrieval-Augmented Generation (RAG), graph analytics, intelligent cyber agents, and other AI-enabled services to produce explanations that remain understandable, verifiable, and directly linked to the underlying semantic knowledge. This capability contributes significantly to trustworthy AI operation and facilitates human oversight of automated cybersecurity analysis. The proposed specification additionally recognises the importance of regulatory accountability. Financial-sector cybersecurity increasingly requires organisations to demonstrate how operational decisions are supported by evidence, how cybersecurity knowledge is governed, and how analytical conclusions can be justified during audits, supervisory reviews, or regulatory investigations. By combining provenance information, lifecycle management, governance attributes, and semantic relationship tracing within a unified audit model, the SKM enables CyberAid to provide structured evidence supporting transparency, accountability, and compliance with evolving regulatory expectations without embedding regulatory logic directly into analytical processes.

Finally, Explainability and Auditability complete the Knowledge Governance Layer of the CyberAid SKM. Together with metadata, provenance, lifecycle management, and governance mechanisms, these principles ensure that semantic knowledge remains not only technically correct but also understandable, trustworthy, reproducible, and accountable throughout its operational lifecycle. This comprehensive governance architecture establishes the SKM as a transparent semantic foundation capable of supporting advanced cybersecurity analytics, trustworthy AI, graph-based reasoning, CTI integration, cyber-risk assessment, and evidence-driven decision support across the CyberAid platform.

4. CTI INTEGRATION FRAMEWORK

Cyber Threat Intelligence provides the continuously updated external awareness required for CyberAid to recognise emerging threats, understand adversarial behaviour, relate technical indicators to known campaigns, and maintain an informed view of the financial-sector threat landscape. Unlike the relatively stable conceptual structures defined by the SKM, CTI is inherently volatile. New indicators appear and expire, vulnerabilities are disclosed and reassessed, attacker infrastructures change, confidence levels evolve, and intelligence reports may be amended as additional evidence becomes available. The value of CTI therefore depends not only on its acquisition, but also on its timely interpretation, quality assessment, contextualisation, and connection with the broader knowledge already maintained by the platform.

Threat intelligence is produced by a wide variety of communities and technical environments. Public repositories, security vendors, vulnerability databases, malware analysis services, research publications, information-sharing communities, operational monitoring systems, financial-sector organisations, and trusted partners may all contribute relevant information. These sources differ considerably in structure, terminology, granularity, update frequency, reliability, access conditions, and permitted use. Some publish structured intelligence through established exchange formats, while others provide semi-structured reports, advisory documents, technical descriptions, or locally generated observations. A direct integration approach based on source-specific data structures would therefore introduce semantic fragmentation, duplicated processing logic, and excessive dependencies between intelligence providers and downstream CyberAid services.

The CTI Integration Framework addresses this challenge by defining the logical flow through which heterogeneous threat information is converted into consistent, traceable, and contextually relevant knowledge. Its role is not limited to collecting external feeds. It establishes the sequence of activities required to acquire intelligence, preserve source information, validate its structure, resolve identifiers, normalise terminology, correlate related objects, assess quality, enrich records with financial-sector context, and align the resulting information with the concepts and associations already defined by the SKM. Through this process, external CTI becomes part of a broader knowledge environment rather than remaining a collection of disconnected indicators or reports.

The framework adopts a clear separation between source acquisition and semantic interpretation. External platforms and connectors remain responsible for obtaining intelligence from their respective providers, whereas the integration layer determines how this information is interpreted and related to CyberAid knowledge. This separation reduces direct coupling between individual CTI sources and analytical services. It also allows sources to be added, replaced, suspended, or restricted without requiring substantial redesign of the knowledge consumers that depend on the resulting intelligence.

Financial-sector contextualisation represents a central requirement of the framework. Generic CTI may describe malicious infrastructure, malware, attack techniques, vulnerabilities, indicators, and adversarial campaigns, but it does not automatically explain their relevance to payment services, banking applications, financial APIs, customer authentication, transaction processing, fraud scenarios, critical business functions, or regulatory responsibilities. The integration process must therefore connect technical intelligence with the financial and organisational domains described in Chapter 3. This contextual enrichment enables threat information to be prioritised according to affected services, operational dependencies, asset criticality, likely business consequences, and applicable resilience or compliance obligations.

The framework also preserves the relational structure of threat intelligence. Individual indicators or vulnerability identifiers provide limited analytical value when detached from their surrounding context. Their significance increases when they remain linked to malware families, campaigns, threat actors, attack procedures, affected products, observed incidents, defensive recommendations, and supporting reports. Accordingly, the integration approach treats CTI objects and their associations as interconnected information that should retain its original evidential and contextual relationships during processing. This is particularly important for graph-oriented analysis and reasoning, where the analytical value derives from traversing connections rather than inspecting isolated records.

Quality and trust considerations are applied throughout the CTI lifecycle rather than appended after ingestion. Intelligence may be incomplete, contradictory, duplicated, obsolete, revoked, or restricted by handling conditions. Each item must therefore retain sufficient information regarding its source, collection time, confidence, reliability, freshness, validation state, external references, and dissemination constraints. These attributes allow downstream services to distinguish verified intelligence from preliminary observations and prevent outdated or insufficiently supported information from being treated as authoritative.

The integration framework is intentionally described at the logical and architectural level. It does not mandate a single CTI platform, connector technology, storage engine, or orchestration mechanism.

OpenCTI, STIX/TAXII-compatible services, vulnerability repositories, graph technologies, and other supporting tools may be used as candidate or validation environments where appropriate, but they do not define the framework itself. The architectural objective is to preserve a stable integration process that remains applicable even if the underlying implementation choices change during subsequent development and integration activities.

Once processed and aligned with the SKM, CTI can support a broad range of CyberAId functions. Graph analytics may use enriched relationships to identify attack paths or suspicious associations. AI components may retrieve evidence-backed context for threat analysis and explanation. Risk-oriented services may connect active campaigns and exploited weaknesses with critical financial operations. Monitoring and incident-response capabilities may compare live observations with known indicators, adversarial patterns, or recommended mitigations. Compliance-oriented services may identify whether an incident or exposure condition has implications for reporting, resilience, or governance responsibilities.

This chapter therefore presents the complete logical path from external intelligence sources to usable CyberAId knowledge. Section 4.1 defines the purpose, responsibilities, and boundaries of the framework. Section 4.2 classifies the intelligence sources that may contribute information, including financial-sector sharing communities and restricted feeds. Section 4.3 describes the processing stages from acquisition through normalisation, correlation, and contextual enrichment. Section 4.4 specifies the STIX-compatible mapping approach used to align CTI objects, associations, and metadata with the SKM. Section 4.5 concludes the chapter by addressing quality assurance, handling restrictions, governance, and controlled bidirectional sharing.

4.1 Purpose and Scope of the CTI Integration Framework

The CTI Integration Framework defines the logical mechanisms through which Cyber Threat Intelligence is systematically incorporated into the CyberAId platform and transformed into structured knowledge suitable for operational cybersecurity analysis. Its primary objective is to provide a consistent integration layer capable of acquiring heterogeneous intelligence, harmonising diverse information structures, preserving semantic meaning, and producing knowledge that can be directly interpreted by the SKM described in Chapter 3. Rather than functioning as an independent threat intelligence platform, the framework acts as the conceptual bridge between external intelligence sources and the knowledge-driven analytical capabilities developed throughout CyberAId. The framework has been designed to support a continuously evolving cyber threat landscape in which intelligence originates from multiple independent providers operating under different technical standards, publication mechanisms, and governance models. Threat intelligence may be distributed through structured standards such as STIX/TAXII, vulnerability repositories, malware analysis platforms, security advisories, information-sharing communities, financial-sector organisations, or internal operational environments. Since these sources employ different data models, vocabularies, update frequencies, confidence levels, and dissemination policies, the integration process must establish a common interpretation layer before this information can be exploited consistently by downstream CyberAId components.

A central responsibility of the CTI Integration Framework is the transformation of heterogeneous threat intelligence into semantically coherent knowledge. This transformation extends beyond syntactic conversion or format normalisation. It involves interpreting the meaning of external cybersecurity information, identifying relationships between independent intelligence objects, preserving contextual information, maintaining provenance, and aligning external observations with the concepts already defined by the SKM. Through this process, threat intelligence evolves from isolated technical artefacts into interconnected knowledge that supports cybersecurity reasoning within the broader financial-sector operational context. The framework also provides the logical foundation for intelligence contextualisation. Generic cyber threat intelligence typically describes adversarial infrastructure, malware families, attack techniques, software vulnerabilities, indicators of compromise, campaigns, and observed behaviours without considering the operational characteristics of specific application domains. Within CyberAId, this information is further interpreted in relation to financial assets, banking services, payment infrastructures,

transaction processing environments, organisational dependencies, business processes, resilience objectives, and regulatory obligations. Contextualisation therefore represents one of the distinguishing characteristics of the framework, allowing technical intelligence to be prioritised according to organisational relevance rather than technical severity alone. Another key objective concerns the preservation of semantic continuity throughout the intelligence integration process. External threat intelligence frequently contains explicit relationships describing associations between adversaries, campaigns, malware, indicators, vulnerabilities, infrastructure, reports, and observed activities. Preserving these associations is essential for enabling graph-based analysis, contextual reasoning, and AI-assisted exploration. The framework therefore maintains not only the individual intelligence objects but also the relationships that connect them, ensuring that the resulting knowledge remains suitable for correlation, dependency analysis, attack path reconstruction, and advanced semantic reasoning. The CTI Integration Framework additionally establishes a logical separation between intelligence acquisition and intelligence utilisation. Activities related to data collection, connector implementation, communication protocols, or platform-specific interfaces are considered implementation concerns that may evolve throughout the development of the CyberAid platform. Conversely, the conceptual integration process remains stable regardless of the technologies employed to retrieve or store threat intelligence. This separation enables the framework to accommodate future intelligence providers, alternative exchange standards, or different implementation technologies without requiring modifications to the underlying integration principles or semantic architecture.

The scope of the framework further encompasses the preservation of essential intelligence characteristics throughout the complete processing pipeline. Information describing source provenance, confidence, reliability, timestamps, handling restrictions, external references, and supporting evidence accompanies threat intelligence as it progresses towards semantic integration. Maintaining these characteristics allows downstream analytical services to evaluate the quality and operational suitability of available intelligence while supporting explainable reasoning, governance, and regulatory accountability.

Although the framework provides the conceptual foundation for integrating cyber threat intelligence, it deliberately excludes implementation-specific decisions regarding software platforms, orchestration technologies, connector development, storage infrastructures, graph databases, or AI processing engines. Likewise, it does not prescribe operational detection logic, incident response procedures, or risk calculation methodologies, which are addressed by other CyberAid components and subsequent development activities. Instead, the framework focuses exclusively on the logical processes required to transform external intelligence into reusable organisational knowledge that can be consumed consistently across the platform. Within the overall CyberAid architecture, the CTI Integration Framework serves as the principal entry point for external cybersecurity knowledge. The integrated intelligence produced by the framework supports multiple downstream capabilities, including semantic graph construction, AI-assisted reasoning, Retrieval-Augmented Generation (RAG), cyber-risk analysis, attack path correlation, security monitoring, operational decision support, resilience assessment, and compliance-oriented services. By providing a unified and semantically consistent representation of threat intelligence, the framework enables these components to operate upon a shared knowledge foundation rather than interacting directly with heterogeneous intelligence sources.

Ultimately, the CTI Integration Framework establishes the conceptual architecture required to convert external cybersecurity intelligence into structured organisational knowledge. By defining a common integration process independent of individual technologies or intelligence providers, it enables CyberAid to maintain a scalable, interoperable, and domain-aware knowledge environment capable of supporting advanced cybersecurity analytics, financial-sector situational awareness, explainable AI, and evidence-driven decision-making.

4.2 CTI Source Categories and ISAC/FS-ISAC Considerations

The effectiveness of the CTI Integration Framework depends largely on the diversity, quality, and complementary nature of the intelligence sources from which cybersecurity knowledge is obtained. No

single provider offers complete visibility across the continuously evolving threat landscape, particularly within the financial sector where cyber threats span technical vulnerabilities, sophisticated fraud campaigns, ransomware operations, supply chain attacks, insider threats, and emerging adversarial techniques. Consequently, CyberAid adopts a multi-source intelligence strategy that combines information originating from different categories of threat intelligence providers, each contributing distinct perspectives, levels of detail, and operational value to the overall knowledge ecosystem. Rather than treating all intelligence sources as equivalent, the framework recognises that different providers contribute different forms of cybersecurity knowledge. Some sources specialise in vulnerability disclosure and software weaknesses, while others focus on adversary behaviour, malware analysis, indicators of compromise, attack techniques, operational incidents, fraud intelligence, or regulatory guidance. Combining these complementary perspectives enables CyberAid to construct a richer representation of the threat landscape than would be possible through any individual intelligence provider. The resulting knowledge is therefore more suitable for semantic correlation, contextual analysis, graph exploration, and AI-assisted reasoning. The identification and governance of CTI sources and sharing communities follows the general information-sharing considerations described in NIST SP 800-150 [15].

A primary category consists of public Cyber Threat Intelligence repositories, which provide openly available information describing malicious infrastructure, indicators of compromise, threat reports, malware campaigns, and emerging cyber threats. These repositories frequently serve as an initial source of operational intelligence and contribute broad situational awareness across multiple industry sectors. Although publicly accessible intelligence offers significant coverage and rapid availability, its quality, completeness, and operational relevance may vary considerably depending on the publisher, update frequency, and validation methodology. Consequently, such information requires additional assessment and contextualisation before becoming part of the SKM.

A second category includes structured threat intelligence platforms and STIX/TAXII-compatible sources. STIX 2.1 defines the structured CTI representation model, while TAXII 2.1 defines an application-layer protocol for exchanging CTI [4], [5]. Together, these standards facilitate automated information exchange, relationship preservation, and interoperability between participating systems. Structured intelligence significantly simplifies semantic integration because entities, relationships, confidence assessments, and supporting metadata are already represented according to recognised exchange models. Nevertheless, differences in implementation practices, optional schema elements, and platform-specific extensions require additional harmonisation before consistent integration can be achieved within the CyberAid knowledge environment. Equally important are vulnerability intelligence sources, which provide detailed information regarding software weaknesses, affected products, exploit availability, remediation guidance, severity assessments, and disclosure timelines. Where available, vulnerability severity is represented using CVSS [9]. Resources such as CVE records [7], the National Vulnerability Database (NVD) [8], vendor security advisories, and vulnerability coordination initiatives contribute essential knowledge describing exploitable conditions that may subsequently be associated with threat actors, attack techniques, affected assets, defensive controls, and organisational risk scenarios. Within CyberAid, vulnerability intelligence therefore represents one of the principal mechanisms through which technical exposure is connected to business and operational context. Another complementary category comprises adversary knowledge repositories, including resources describing attacker tactics, techniques, procedures, operational behaviours, and campaign characteristics. MITRE ATT&CK [6] provides structured descriptions of adversarial behaviour that enable CyberAid to interpret observed activities according to recognised attack methodologies rather than isolated technical events. This behavioural perspective strengthens semantic correlation by connecting technical indicators with higher-level attack strategies, facilitating attack path analysis, adversary profiling, and contextual threat assessment.

The framework also considers malware intelligence and security research as important contributors to cybersecurity knowledge. Malware analysis reports, reverse engineering studies, exploit analyses, proof-of-concept publications, and vendor research frequently provide behavioural observations that complement structured threat intelligence. Although these resources may vary considerably in format and structure, they often contain valuable contextual information regarding attack evolution, malware

capabilities, command-and-control infrastructure, exploitation techniques, or defensive recommendations. Such information enriches the SKM by extending the semantic context surrounding malicious activities beyond standardised threat intelligence objects. Beyond traditional cybersecurity repositories, CyberAid also incorporates regulatory and guidance sources that influence cybersecurity governance within the financial sector. Publications issued by European institutions, supervisory authorities, standards organisations, national cybersecurity agencies, and sector-specific regulators provide valuable information regarding resilience expectations, reporting obligations, security controls, and emerging regulatory developments. While these sources do not constitute cyber threat intelligence in the conventional sense, they provide essential contextual knowledge that enables CyberAid to relate technical cybersecurity observations to applicable governance and compliance requirements.

A further source category includes internal organisational intelligence generated during operational cybersecurity activities. Security monitoring systems, incident investigations, forensic analyses, vulnerability assessments, penetration testing, digital twin simulations, and historical organisational experience all contribute valuable knowledge describing the operational environment in which CyberAid is deployed. Integrating internally generated intelligence alongside external threat information enables the platform to correlate global threat developments with locally observed behaviours, thereby improving situational awareness and strengthening organisation-specific risk assessment. Particular significance is attributed to Information Sharing and Analysis Centres (ISACs) and, more specifically, to the Financial Services Information Sharing and Analysis Center (FS-ISAC) [14]. These collaborative communities facilitate the exchange of sector-specific intelligence among trusted participants, often providing timely information regarding emerging campaigns, fraud trends, targeted attacks, operational vulnerabilities, and defensive practices affecting financial institutions. Their value lies not only in the technical intelligence they disseminate but also in the operational context, sector-specific expertise, and collaborative situational awareness they provide. However, participation in such communities is generally governed by membership agreements, licensing conditions, confidentiality obligations, and information handling rules. Consequently, the CTI Integration Framework considers ISAC-derived intelligence as a highly valuable but conditionally available source whose integration depends on organisational participation and compliance with applicable sharing policies. Because intelligence sources differ substantially in quality and operational characteristics, CyberAid does not prioritise any individual provider as a definitive source of truth. Instead, intelligence originating from multiple repositories is evaluated collectively, allowing corroboration across independent observations, identification of inconsistencies, and enrichment through complementary evidence. This multi-source approach increases analytical robustness while reducing dependence on individual providers whose coverage, availability, or publication practices may change over time. It also aligns with the broader objective of constructing a resilient and adaptable knowledge ecosystem capable of evolving alongside the cybersecurity landscape.

The integration framework intentionally remains independent of specific commercial products or intelligence providers. Platforms capable of supporting structured CTI exchange, such as OpenCTI [13], may be employed as candidate environments for experimentation, inspection, or validation during research and development activities, while established standards such as STIX [4] and TAXII [5] facilitate interoperability wherever applicable. Nevertheless, the architectural principles described in this deliverable remain applicable regardless of the technologies ultimately adopted for operational deployment, ensuring that the framework preserves its conceptual validity as CyberAid continues to evolve.

Table 9 - Representative CTI Source Categories

Source Category	Representative Sources	Primary Contribution	Typical Data
Public CTI Repositories	Open CTI feeds, CERT advisories, threat reports	General threat awareness	IoCs, campaigns, reports

Source Category	Representative Sources	Primary Contribution	Typical Data
Structured CTI Platforms	STIX/TAXII-compatible platforms, OpenCTI (reference environment)	Machine-readable CTI	Threat objects, relationships, metadata
Vulnerability Intelligence	CVE, NVD, vendor advisories	Software exposure intelligence	Vulnerabilities, CVSS, affected products
Adversary Knowledge	MITRE ATT&CK	Behavioural intelligence	Tactics, techniques, procedures
Malware Intelligence	Security vendors, malware research	Malware behaviour	Families, capabilities, infrastructure
Regulatory & Guidance Sources	ENISA, NIST, DORA guidance, OWASP	Governance context	Controls, recommendations, regulatory guidance
Internal Organisational Intelligence	SOCs, vulnerability assessments, incident investigations	Organisation-specific intelligence	Events, findings, incidents
ISAC / FS-ISAC Communities	Financial sector information-sharing communities	Sector-specific situational awareness	Alerts, fraud trends, coordinated intelligence

4.3 CTI Ingestion, Normalization and Correlation

The CyberAid CTI Integration Framework establishes a logical processing pipeline through which heterogeneous cyber threat intelligence is systematically transformed into structured knowledge suitable for semantic representation within the SKM. The objective of this pipeline is not merely to collect threat information from multiple providers but to ensure that intelligence originating from diverse technical environments is interpreted consistently, enriched with additional context, and integrated into a unified knowledge space that supports graph analytics, AI-assisted reasoning, cyber-risk analysis, and operational decision support.

Threat intelligence is characterised by considerable diversity in terms of structure, semantics, publication mechanisms, and operational maturity. Information describing adversaries, attack techniques, software vulnerabilities, malware families, indicators of compromise, campaigns, reports, and defensive recommendations may be distributed through structured exchange standards, semi-structured documents, vendor advisories, vulnerability repositories, collaborative information-sharing communities, or internally generated operational observations. Without a common integration process, these heterogeneous sources would remain isolated, preventing effective knowledge correlation and limiting the analytical capabilities of the CyberAid platform. To address this challenge, the CTI Integration Framework defines a sequence of logical processing stages that progressively transform raw intelligence into semantically enriched

knowledge. Although individual implementation technologies may vary, the conceptual flow remains stable throughout the platform lifecycle and provides a common integration model for all participating intelligence sources.

The first stage of the pipeline is intelligence acquisition, during which relevant cyber threat information is collected from authorised external and internal providers. Acquisition encompasses both periodically retrieved and event-driven intelligence, depending on the characteristics of the originating source. At this stage, the objective is to preserve the original structure of the received information while recording essential contextual information, including source identifiers, acquisition timestamps, access restrictions, handling instructions, and accompanying metadata. No semantic interpretation is performed during acquisition; instead, the framework focuses on ensuring that incoming intelligence is captured in a complete and traceable manner. Following acquisition, the intelligence progresses to the validation and pre-processing stage. Incoming information is examined to verify structural completeness, identify malformed or incomplete records, detect duplicated observations, and ensure compatibility with the supported integration process. Where necessary, preliminary transformations may be applied to resolve syntax inconsistencies, character encoding issues, missing mandatory attributes, or unsupported data representations. Validation at this stage aims to establish a consistent baseline before semantic processing begins while preserving the integrity of the original intelligence. The next processing stage concerns normalisation, through which heterogeneous cybersecurity information is translated into a common conceptual representation. Different intelligence providers frequently employ alternative naming conventions, classification schemes, identifiers, and structural representations for similar cybersecurity concepts. The normalisation process harmonises these differences by aligning equivalent concepts, resolving identifier inconsistencies, standardising terminology, and preserving semantic equivalence across independent intelligence sources. Rather than modifying the meaning of the original intelligence, normalisation establishes a consistent representation that enables subsequent correlation and semantic interpretation.

Once a common representation has been established, the framework performs semantic correlation across the available intelligence. Individual cybersecurity observations rarely provide sufficient analytical value when considered independently. Correlation therefore identifies semantic associations between threat actors, attack techniques, vulnerabilities, malware families, campaigns, indicators, observed incidents, defensive controls, and organisational assets. Information originating from different providers may complement, confirm, or extend existing knowledge, allowing multiple observations to contribute to a richer and more comprehensive understanding of the evolving threat landscape. This correlation process significantly enhances situational awareness while reducing fragmentation across heterogeneous intelligence repositories. The correlated intelligence subsequently undergoes contextual enrichment, during which technical cybersecurity observations are associated with the operational environment of the financial sector. Generic CTI describing adversarial infrastructure or exploitation activity is supplemented with organisational context, including affected financial assets, payment infrastructures, banking services, transaction processing systems, business processes, resilience objectives, and applicable regulatory considerations. This enrichment process enables CyberAid to interpret identical technical observations differently depending on the organisational environment in which they occur, thereby improving prioritisation and supporting risk-informed cybersecurity analysis. Following enrichment, the framework performs semantic integration with the SKM. At this stage, the processed intelligence is aligned with the concepts, relationships, metadata structures, and governance principles defined in Chapter 3. Rather than preserving intelligence as isolated source-specific records, the framework represents the information through the common semantic abstractions of the SKM, enabling interoperability across multiple intelligence providers while maintaining provenance, contextual associations, confidence information, and supporting evidence. This integration transforms external cyber threat intelligence into reusable organisational knowledge that can be exploited consistently throughout the CyberAid platform.

Throughout the complete processing pipeline, the framework preserves the contextual characteristics accompanying the original intelligence. Provenance information, confidence assessments, timestamps, external references, handling restrictions, confidence indicators, TLP classifications, source reliability, and

validation status remain associated with each knowledge element as it progresses through the different processing stages. Preserving these contextual characteristics ensures that downstream analytical components can evaluate the quality and trustworthiness of the available intelligence while supporting explainable reasoning and governance requirements. An equally important characteristic of the processing pipeline is its support for incremental knowledge evolution. Newly acquired intelligence is not treated as an isolated dataset replacing existing knowledge but as additional evidence that may confirm, refine, extend, or supersede previous observations. Existing semantic entities may therefore be enriched with newly available intelligence, relationships may be expanded as additional contextual information becomes available, and confidence assessments may evolve as corroborating evidence accumulates from multiple independent sources. This incremental approach enables the SKM to function as a continuously evolving knowledge ecosystem while preserving historical continuity and semantic consistency. The logical pipeline defined by the CTI Integration Framework remains independent of specific implementation technologies, orchestration mechanisms, or processing platforms. Whether intelligence is exchanged through STIX/TAXII services, inspected using candidate environments such as OpenCTI, processed through graph-oriented infrastructures, or integrated using future CyberAId components, the conceptual stages of acquisition, validation, normalisation, correlation, contextual enrichment, and semantic integration remain unchanged. This separation between logical architecture and implementation technologies provides the flexibility required for long-term platform evolution while preserving the conceptual consistency of the intelligence integration process.

The output of this pipeline is a coherent body of semantically enriched cybersecurity knowledge that can be consumed consistently by the various analytical capabilities of CyberAId. Graph analytics, intelligent cyber agents, Retrieval-Augmented Generation (RAG), cyber-risk assessment services, monitoring components, digital twin environments, incident response mechanisms, and compliance-oriented functions all operate upon the same integrated knowledge foundation. By transforming fragmented external intelligence into structured and contextualised organisational knowledge, the CTI Integration Framework establishes the information backbone that enables CyberAId to deliver advanced cybersecurity capabilities across the financial sector.




Figure 2 - Recommended Logical Processing Flow

4.4 STIX-Compatible CTI-to-SKM Mapping Approach

The CTI Integration Framework transforms externally acquired cyber threat intelligence into semantically consistent knowledge by aligning standardised CTI representations with the concepts defined by the CyberAId SKM. As a significant proportion of structured cyber threat intelligence is exchanged using the Structured Threat Information Expression (STIX) standard, the integration process adopts a STIX-compatible mapping approach that enables externally published intelligence to be interpreted within the semantic architecture introduced in Chapter 3.

The mapping approach establishes a conceptual correspondence between CTI entities and the semantic concepts maintained within the SKM. Standardised intelligence objects describing threat actors, malware, attack techniques, campaigns, vulnerabilities, indicators, reports, identities, infrastructure, observed data, and relationships are interpreted according to their semantic role before being associated with the appropriate SKM entities. This abstraction allows CyberAid to remain interoperable with widely adopted CTI standards while avoiding direct dependence on a particular exchange model or implementation technology. A fundamental design principle of the mapping process is semantic preservation. Intelligence objects should retain their original meaning, contextual information, and evidential value throughout the transformation process. Consequently, the mapping does not simply rename STIX objects or translate individual attributes into alternative formats. Instead, it interprets the operational significance of each intelligence element and associates it with the corresponding semantic structures already defined by the SKM. This ensures that downstream analytical services operate upon a coherent conceptual model rather than a heterogeneous collection of source-specific representations.

The mapping process distinguishes between three complementary dimensions of semantic alignment: entity mapping, relationship mapping, and metadata mapping. Together, these dimensions enable CyberAid to preserve both the informational content and the contextual characteristics of externally acquired intelligence while maintaining interoperability with established CTI standards.

Entity mapping focuses on the correspondence between individual intelligence objects and SKM concepts. Objects describing adversaries, malicious software, vulnerabilities, indicators, campaigns, defensive measures, reports, identities, infrastructure, and observed activities are associated with the equivalent semantic entities represented within the SKM. During this process, intelligence originating from different providers may converge towards the same conceptual representation, enabling consistent interpretation despite variations in naming conventions, publication formats, or provider-specific modelling practices. The resulting abstraction reduces semantic duplication while facilitating knowledge reuse across multiple analytical services. Equally important is relationship mapping, through which the associations originally present within threat intelligence are preserved inside the SKM. Threat actors may employ specific attack techniques, malware may exploit known vulnerabilities, campaigns may utilise malicious infrastructure, indicators may be linked to observed activities, and defensive controls may mitigate particular attack scenarios. These semantic associations are maintained throughout the mapping process to ensure that graph analytics, contextual reasoning, dependency analysis, and AI-assisted exploration can exploit the complete relational structure of the available intelligence. Preserving relationships is particularly important because the analytical value of cyber threat intelligence often derives from the connections between entities rather than from the entities themselves. The third mapping dimension concerns metadata alignment. Structured threat intelligence typically includes descriptive information regarding confidence, provenance, timestamps, external references, handling restrictions, source reliability, version information, and other contextual attributes. Rather than treating these elements as secondary information, the mapping process associates them with the metadata, provenance, lifecycle, and governance structures introduced in Section 3.6. This ensures that externally acquired intelligence preserves its operational characteristics after integration while enabling CyberAid to support traceability, explainability, knowledge governance, and evidence-based reasoning.

Although the mapping process is designed to support STIX-compatible intelligence, it is intentionally model-oriented rather than format-dependent. Other structured or semi-structured intelligence sources may be incorporated using equivalent semantic alignment principles provided that their concepts, relationships, and contextual information can be associated with the abstractions defined by the SKM. This design choice reduces dependency on individual standards while preserving interoperability with future intelligence exchange mechanisms that may emerge throughout the lifecycle of the CyberAid platform.

The mapping process also accommodates semantic enrichment. External CTI frequently provides generic descriptions of cyber threats without organisation-specific interpretation. Once intelligence objects have been associated with SKM concepts, they may be complemented with additional contextual information describing affected financial services, critical business assets, payment infrastructures, transaction

processes, organisational dependencies, resilience objectives, regulatory obligations, or historical operational observations. Consequently, the resulting knowledge extends beyond the information originally published by the intelligence provider, enabling CyberAId to support richer analysis tailored to the financial sector. The proposed mapping approach further supports semantic consistency across heterogeneous intelligence providers. Multiple sources may describe the same adversary, campaign, vulnerability, or malware family using different identifiers, naming conventions, or publication styles. Through semantic alignment, these independent observations may be associated with common SKM concepts while preserving their individual provenance, confidence, and supporting evidence. This capability improves knowledge consolidation, reduces unnecessary duplication, and strengthens multi-source intelligence correlation without sacrificing source traceability. Because the mapping process forms the interface between external CTI and the internal knowledge architecture, transparency remains an essential design objective. Every semantic correspondence should remain traceable to its originating intelligence object, allowing analysts to understand how externally published information has been interpreted and represented within the SKM. This traceability supports explainable AI, facilitates validation of mapping decisions, and enables future refinement as intelligence standards and cybersecurity practices continue to evolve. Structured CTI inspection and preliminary mapping experiments were performed using OpenCTI as a representative reference environment for exploring STIX-compatible cyber threat intelligence. These activities supported the validation of the conceptual mapping approach described in this deliverable and provided practical insight into the semantic correspondence between externally published intelligence and the concepts defined within the SKM. The use of OpenCTI should therefore be regarded as a validation-oriented environment rather than a prescribed implementation component of the CyberAId architecture. Thus, the STIX-compatible mapping approach provides the semantic bridge through which externally acquired cyber threat intelligence becomes integrated organisational knowledge. By preserving entities, relationships, metadata, provenance, and contextual information while aligning them with the abstractions of the SKM, the mapping process enables CyberAId to maintain interoperability with established CTI standards without compromising the consistency or extensibility of its internal knowledge architecture.

Table 10 - Representative STIX-to-SKM Mapping

STIX Concept	Corresponding SKM Concept	Mapping Purpose
Threat Actor	Threat Actor	Adversary representation
Attack Pattern	Attack Technique	Behaviour modelling
Malware	Malware	Malicious software knowledge
Vulnerability	Vulnerability	Exposure representation
Campaign	Campaign	Coordinated attack activity
Indicator	Indicator	Detection knowledge
Infrastructure	Infrastructure	Supporting attack resources
Identity	Organisation / Identity	Contextual entities

STIX Concept	Corresponding SKM Concept	Mapping Purpose
Observed Data	Observation	Operational evidence
Report	CTI Report	Intelligence documentation
Relationship	Semantic Relationship	Knowledge graph connectivity
Marking Definition	Governance Metadata	Handling and dissemination rules

4.5 Data Quality, Governance and Bidirectional Sharing

The value of cyber threat intelligence depends not only on the amount of information collected but also on its quality, reliability, traceability, and suitability for operational use. Threat intelligence frequently originates from multiple independent providers operating under different publication methodologies, validation practices, confidence models, and information-sharing policies. Consequently, the CTI Integration Framework incorporates a comprehensive set of quality assessment and governance principles to ensure that integrated intelligence remains accurate, trustworthy, consistent, and suitable for supporting cybersecurity operations within the CyberAid platform. Quality assurance begins before intelligence is incorporated into the SKM. Every intelligence element should be evaluated according to a combination of structural, semantic, and operational criteria that collectively determine its suitability for integration. Structural validation confirms that the information can be interpreted correctly by the integration pipeline, while semantic validation examines whether the intelligence can be associated with existing concepts, relationships, and contextual information without introducing inconsistencies or ambiguity. Operational validation considers factors such as timeliness, completeness, relevance, and practical usefulness within the financial-sector cybersecurity context.

A fundamental quality dimension concerns source reliability. Intelligence providers differ significantly in their expertise, collection methodologies, publication frequency, and operational credibility. Rather than assuming uniform trust across all providers, the framework preserves information describing the reliability of the originating source, enabling downstream analytical services to distinguish highly trusted intelligence from preliminary observations or unverified reports. This distinction supports evidence-based reasoning and prevents analytical conclusions from relying exclusively on intelligence of uncertain quality.

Closely related to source reliability is the assessment of confidence associated with individual intelligence elements. Confidence reflects the degree of certainty assigned to a specific observation, report, or analytical conclusion and may evolve as additional evidence becomes available. Intelligence collected from independent providers may reinforce or challenge previous assessments, allowing confidence levels to be refined throughout the operational lifecycle of the knowledge. By preserving confidence information alongside semantic entities and relationships, CyberAid enables analysts and automated services to interpret intelligence according to its evidential strength rather than treating all observations as equally authoritative. Another important quality consideration concerns timeliness and freshness. Cyber threats evolve rapidly, and intelligence that was operationally valuable at one point may become obsolete as attacker infrastructure changes, vulnerabilities are remediated, malware variants evolve, or campaigns conclude. The CTI Integration Framework therefore preserves temporal information describing publication dates, observation times, update history, and review status, enabling CyberAid to differentiate current operational intelligence from historical information retained primarily for analytical or forensic purposes. This temporal awareness contributes to more accurate threat assessment and reduces the risk of relying on outdated intelligence during operational decision-making. The framework also incorporates mechanisms for deduplication and knowledge consolidation. Because multiple intelligence providers frequently report

the same adversary, campaign, vulnerability, or indicator, repeated integration of equivalent information may lead to unnecessary redundancy and fragmented knowledge representations. Rather than maintaining isolated copies of semantically equivalent observations, the framework supports the consolidation of related intelligence while preserving the provenance, confidence, and contextual information contributed by each independent source. This approach reduces duplication without compromising evidential traceability and strengthens the completeness of the resulting knowledge.

A further quality aspect concerns the management of revocation and superseded intelligence. Cybersecurity information occasionally becomes invalid due to corrected analyses, false positive detections, infrastructure changes, withdrawn advisories, or revised threat assessments. The integration process therefore accommodates mechanisms that identify obsolete or superseded intelligence while preserving its historical significance. Instead of permanently removing such information, the framework maintains appropriate lifecycle status and provenance records, allowing historical analyses to remain reproducible while ensuring that operational services rely primarily on current and valid intelligence. Beyond quality assessment, the CTI Integration Framework establishes governance principles governing the handling of integrated intelligence. Information acquired from external providers frequently carries dissemination restrictions, confidentiality obligations, licensing conditions, or contractual limitations that determine how it may be stored, processed, or redistributed. Consequently, governance attributes accompany intelligence throughout the integration process, ensuring that handling policies remain associated with the corresponding semantic knowledge after integration into the SKM. Particular attention is given to Traffic Light Protocol (TLP) markings and equivalent information handling classifications. Many intelligence providers use TLP or comparable dissemination schemes to communicate the permitted audience and redistribution constraints associated with shared information. Rather than treating these markings as external administrative metadata, the CTI Integration Framework preserves them as operational governance attributes that influence subsequent sharing and consumption activities. This approach enables CyberAId to support collaborative intelligence utilisation while respecting the obligations defined by the originating information provider. The framework also adopts privacy-by-design and confidentiality-by-design principles during intelligence processing. Although cyber threat intelligence primarily concerns technical observations, operational environments may occasionally generate information containing organisationally sensitive data, incident-specific details, or references to identifiable entities. The integration process therefore considers data minimisation, controlled visibility, proportional information disclosure, and appropriate handling restrictions wherever applicable, ensuring alignment with organisational governance requirements and relevant European regulatory frameworks without reducing the analytical value of the resulting knowledge.

Another essential governance objective is the preservation of integrity and auditability throughout the intelligence lifecycle. Every significant modification affecting integrated intelligence—including validation, enrichment, confidence updates, semantic mapping, governance decisions, or lifecycle transitions—should remain traceable through appropriate provenance and governance information. This capability enables CyberAId to reconstruct the evolution of integrated intelligence, justify analytical conclusions, support forensic investigations, and demonstrate compliance with organisational policies and regulatory expectations. The CTI Integration Framework additionally supports the concept of controlled bidirectional intelligence sharing. While the primary objective of the framework is the acquisition and integration of external threat intelligence, CyberAId may also generate enriched knowledge, analytical findings, contextual correlations, or validated observations that could provide value to trusted external communities. Subject to organisational approval, licensing conditions, confidentiality obligations, and applicable information-sharing agreements, selected knowledge may therefore be disseminated back to authorised partners, collaborative intelligence communities, or sector-specific information-sharing initiatives. Such exchanges should preserve provenance, governance attributes, confidence information, and applicable dissemination restrictions, ensuring that externally shared knowledge remains both traceable and responsibly governed. The bidirectional sharing capability should not be interpreted as unrestricted intelligence publication. Rather, it represents a governance-controlled mechanism through which validated knowledge may contribute to collective cyber resilience while respecting ownership, confidentiality,

intellectual property rights, and information handling obligations. This collaborative perspective aligns with the broader objective of strengthening situational awareness across trusted cybersecurity communities without compromising organisational security or regulatory compliance.

Finally, the principles presented throughout this section ensure that integrated cyber threat intelligence remains not only technically usable but also operationally trustworthy, legally manageable, and organisationally accountable. By combining quality assessment, provenance preservation, governance policies, handling restrictions, lifecycle awareness, and controlled information sharing, the CTI Integration Framework establishes a resilient foundation for maintaining reliable cybersecurity knowledge throughout the CyberAid platform. Together with the SKM defined in Chapter 3 and the integration processes described throughout this chapter, these governance mechanisms complete the conceptual architecture supporting intelligent, explainable, and collaborative cyber threat intelligence management within CyberAid.

5. KNOWLEDGE REPRESENTATION AND CONSUMPTION

The previous chapters established the conceptual foundations of the CyberAid SKM and described the mechanisms through which heterogeneous cyber threat intelligence is acquired, transformed, and semantically integrated into a unified knowledge environment. Having defined both the semantic architecture and the intelligence integration process, the next logical step is to examine how this knowledge can be represented computationally and subsequently exploited by the various functional components of the CyberAid platform. Knowledge representation plays a fundamental role in ensuring that cybersecurity information remains machine-interpretable, semantically consistent, and suitable for advanced analytical processing. A well-defined representation model enables cybersecurity concepts, relationships, metadata, and contextual information to be organised in a structured manner while preserving interoperability across heterogeneous data sources and analytical services. Rather than prescribing a particular implementation technology, the SKM defines an abstract semantic representation that may be realised through different computational paradigms depending on deployment requirements, scalability considerations, and future technological evolution. The representation approach proposed within CyberAid is designed to support multiple complementary objectives. At the semantic level, it enables the formal description of cybersecurity entities, relationships, organisational context, financial-domain concepts, and governance information. At the analytical level, it facilitates graph-based exploration, semantic querying, contextual reasoning, and multi-domain knowledge correlation. At the intelligence level, it provides the structured foundation required for AI-assisted analysis, Retrieval-Augmented Generation (RAG), explainable decision support, and future agent-oriented cybersecurity services. These complementary capabilities allow the SKM to function as the common semantic foundation upon which diverse CyberAid services can consistently operate.

This chapter therefore focuses on the computational representation of the SKM and its utilisation throughout the CyberAid architecture. Section 5.1 introduces the principal representation models through which the conceptual knowledge defined in previous chapters may be expressed independently of specific implementation technologies. Section 5.2 examines the semantic characteristics that make the SKM suitable for AI-driven processing and large language model integration, while Section 5.3 describes how the resulting knowledge is consumed by the different analytical and operational components of the CyberAid ecosystem, providing a unified knowledge foundation for cybersecurity intelligence, financial risk assessment, monitoring, compliance, and decision support services.

5.1 Knowledge Representation Models

The SKM defined throughout this deliverable represents the conceptual foundation upon which CyberAid organises cybersecurity knowledge, cyber threat intelligence, financial context, organisational information, and governance attributes. While the SKM is inherently a conceptual semantic model, its practical utilisation requires computational representations capable of expressing cybersecurity entities, relationships, contextual information, and associated metadata in a structured, machine-interpretable form.

Consequently, the CyberAId architecture adopts a representation-oriented perspective that separates the conceptual definition of knowledge from the computational models used to express it, thereby preserving implementation independence while enabling interoperability across diverse analytical environments. This distinction between the conceptual model and its computational representations constitutes an important architectural principle. The conceptual model defines what knowledge exists, how cybersecurity concepts relate to one another, and which semantic constraints govern these relationships. Computational representations, by contrast, define how this conceptual knowledge may be expressed, organised, queried, and exploited by software components. Maintaining this separation ensures that the SKM remains stable even as implementation technologies evolve, allowing different representation mechanisms to coexist without altering the underlying semantic architecture.

At the highest level of abstraction, the SKM may be represented through conceptual taxonomies that organise cybersecurity knowledge into coherent semantic domains. These taxonomies classify concepts according to their functional role, distinguishing, for example, threat actors, attack techniques, vulnerabilities, malware, defensive controls, organisational assets, financial processes, regulatory concepts, and governance information. Such hierarchical organisation establishes semantic consistency while facilitating navigation across complex cybersecurity knowledge spaces. Although taxonomies provide structural organisation, they primarily describe conceptual classification rather than the richer semantic relationships required for advanced analytical reasoning. To capture these richer semantics, the conceptual model may be expressed through ontology-oriented representations. Ontologies extend hierarchical classification by defining semantic classes together with the formal relationships that connect them. Within this representation paradigm, cybersecurity concepts become explicitly defined entities whose characteristics, associations, and constraints can be described in a structured and logically consistent manner. Rather than focusing solely on terminology, ontology-oriented representations establish an explicit semantic vocabulary that enables different CyberAId components to interpret cybersecurity knowledge according to a common conceptual understanding. Within such representations, semantic classes define the principal categories of knowledge maintained by the SKM. These classes describe the major conceptual entities identified throughout Chapters 3 and 4, including adversaries, attack techniques, vulnerabilities, indicators, campaigns, malware, incidents, assets, business processes, defensive mechanisms, regulatory concepts, and governance information. Each class encapsulates the semantic characteristics shared by a particular category of cybersecurity knowledge while remaining sufficiently extensible to accommodate future concepts introduced during the evolution of the CyberAId platform.

The semantic description of these classes is further enriched through object properties and data properties, which capture complementary aspects of cybersecurity knowledge. Object properties express semantic relationships between conceptual entities, allowing the representation model to describe, for example, how threat actors employ attack techniques, how vulnerabilities affect assets, how defensive controls mitigate risks, or how incidents relate to organisational processes. Data properties, in contrast, associate semantic entities with descriptive information such as identifiers, confidence values, timestamps, severity assessments, lifecycle states, or governance attributes. Together, these complementary property types provide the expressive capability required to represent both the structural organisation and the descriptive characteristics of integrated cybersecurity knowledge. Beyond individual entities and their attributes, the SKM emphasises the representation of semantic associations as first-class architectural elements. Cybersecurity knowledge derives much of its analytical value from the relationships connecting different concepts rather than from isolated entities considered independently. Consequently, semantic associations describe behavioural dependencies, contextual interactions, causal relationships, organisational impacts, temporal evolution, and governance constraints across multiple cybersecurity domains. Representing these associations explicitly enables CyberAId to support richer contextual analysis, explainable reasoning, and multidimensional knowledge exploration. While ontology-oriented representations provide a strong semantic foundation, many cybersecurity analysis scenarios also benefit from graph-oriented representations. Graph structures naturally express interconnected cybersecurity knowledge by representing semantic entities as nodes and their associations as edges within a connected knowledge structure. This representation facilitates intuitive modelling of complex attack scenarios, infrastructure

dependencies, organisational relationships, vulnerability propagation, incident evolution, and cross-domain intelligence correlations. Because cybersecurity ecosystems are inherently relational, graph-oriented representations provide an effective computational abstraction for navigating highly interconnected knowledge environments.

Within graph representations, nodes correspond to the conceptual entities defined by the SKM, while edges represent the semantic relationships connecting these entities. As knowledge evolves through continuous intelligence integration, the resulting graph forms increasingly complex paths that capture sequences of related cybersecurity events, adversarial behaviours, attack chains, defensive responses, organisational dependencies, or financial impacts. These interconnected structures enable analysts and automated services to explore cybersecurity knowledge beyond isolated observations, revealing indirect relationships and emerging patterns that may otherwise remain difficult to identify. The representation model also supports relationship-oriented querying, allowing knowledge consumers to retrieve information based not only on individual entities but also on the semantic connections that exist between them. Rather than searching for isolated cybersecurity objects, analytical components may traverse relationships, follow behavioural dependencies, investigate attack paths, identify affected organisational assets, analyse vulnerability propagation, or explore regulatory implications through connected semantic structures. This relationship-centric perspective significantly enhances the analytical capabilities available to CyberAid services while remaining independent of any particular query language or graph processing technology. An equally important characteristic of the proposed representation approach is its support for multiple complementary representation models operating over the same conceptual foundation. The SKM is not tied exclusively to taxonomies, ontologies, or graph structures; instead, these representations provide alternative computational views of the same underlying semantic architecture. Taxonomies emphasise conceptual organisation, ontology-oriented representations formalise semantic meaning, and graph structures facilitate relationship-centric analysis. Together, they provide complementary perspectives that collectively maximise interoperability, analytical flexibility, and future extensibility without altering the conceptual integrity of the SKM itself.

Thus, the representation models presented in this section provide the computational foundation through which the semantic architecture specified throughout this deliverable becomes suitable for practical analytical exploitation. By separating conceptual knowledge from implementation technologies while supporting hierarchical classification, semantic modelling, graph-oriented representation, and relationship-based exploration, the SKM establishes a flexible and technology-independent representation framework capable of supporting the diverse analytical requirements of the CyberAid platform throughout its lifecycle.

5.2 Semantic and LLM-Ready Representation

The computational representation models described in the previous section establish the structural foundation of the SKM. However, effective cybersecurity knowledge representation extends beyond the ability to organise information into semantic entities and relationships. Modern cybersecurity platforms increasingly require knowledge structures that support intelligent analysis, contextual retrieval, automated reasoning, explainable decision-making, and collaboration between human analysts and AI-based services. Consequently, the SKM is designed not only as a semantic representation of cybersecurity knowledge but also as an AI-ready knowledge foundation capable of supporting diverse analytical paradigms across the CyberAid platform. The concept of an AI-ready representation refers to the ability of structured knowledge to be interpreted, queried, enriched, and exploited by intelligent computational services without requiring continuous manual restructuring or domain-specific transformations. Rather than treating cybersecurity information as isolated documents or independent datasets, the SKM organises knowledge into interconnected semantic structures that preserve context, relationships, provenance, governance information, and domain semantics. This representation enables AI-based services to interpret cybersecurity observations within their broader operational context instead of relying exclusively on statistical associations or textual similarity.

For semantic indexing and retrieval, the SKM distinguishes four logical retrievable-unit types: a source document or document section, a semantic entity, a typed relationship, and a bounded subgraph. Each retrievable unit shall be independently addressable and shall preserve a reference to its authoritative source. A model-ready representation shall contain a concise textual rendering of the selected unit together with structured metadata. The minimum indicative retrieval-result representation comprises `element_id`, `element_type`, `title_or_label`, `content`, `source_reference`, `provenance_reference`, `classification`, `handling_constraints`, `confidence`, `valid_from`, `valid_until`, `knowledge_partition`, `semantic_domain` and, where applicable, `graph_scope`. A bounded subgraph shall additionally identify its included nodes, relationships and traversal boundary. The semantic specification allows candidate knowledge to be filtered by semantic domain, partition, graph scope, element type, temporal validity, source, classification and confidence. It also permits a consumer to request a maximum result count. The retrieval API, embedding model, tokenisation method, ranking algorithm, query execution and access-control enforcement are defined by the corresponding implementation and orchestration components.

A fundamental capability supported by the representation model is semantic retrieval. Traditional keyword-based retrieval mechanisms frequently struggle to capture the contextual relationships that exist between cybersecurity concepts. By contrast, semantic retrieval exploits the explicit relationships represented within the SKM to identify knowledge that is conceptually relevant even when identical terminology is not used. Queries may therefore retrieve information based on semantic similarity, behavioural relationships, organisational context, or shared dependencies rather than simple lexical matching, providing analysts with richer and more meaningful access to integrated cybersecurity knowledge. The representation also establishes the foundation for graph analytics, where cybersecurity knowledge is explored through the interconnected relationships linking entities across multiple semantic domains. Instead of analysing individual observations independently, graph-oriented analytical techniques examine the structural organisation of the knowledge itself, enabling the identification of attack paths, infrastructure dependencies, adversarial behaviour patterns, vulnerability propagation chains, organisational impacts, and complex multi-stage attack scenarios. Because these analyses exploit explicit semantic relationships, they provide deeper contextual insight than isolated event analysis while supporting explainable analytical outcomes. An equally important capability concerns Retrieval-Augmented Generation (RAG). Large language models possess powerful reasoning and language generation capabilities but may produce incomplete, outdated, or unsupported responses when operating without access to authoritative domain knowledge. The SKM addresses this limitation by providing a structured semantic knowledge source that can be consulted during the retrieval phase preceding response generation. Rather than relying exclusively on the internal parameters of a language model, Retrieval-Augmented Generation enables AI services to ground their responses in current, traceable, and contextually relevant cybersecurity knowledge maintained within the CyberAId knowledge environment. This approach improves factual consistency, supports evidence-based reasoning, and enables generated outputs to remain aligned with the continuously evolving cybersecurity landscape. In this context, RAG-enabled services are not assumed to retrieve unrestrictedly from the complete CyberAId knowledge environment. Retrieval may instead operate over one or more authorised knowledge partitions, with the eligible knowledge scope determined at retrieval time according to the security context of the request and the applicable governance policies defined by the platform. This approach supports least-privilege knowledge access while maintaining a common semantic representation across independently authorised knowledge domains. The SKM provides the partition, provenance, classification and handling metadata required to support this process. Query execution, Security Context evaluation, result filtering and enforcement of the permitted retrieval scope belong to the consuming retrieval and orchestration components. Beyond retrieval, the semantic representation supports AI-assisted reasoning by preserving explicit relationships, contextual dependencies, and governance information that enable intelligent services to interpret cybersecurity situations in a more structured manner. Rather than analysing isolated indicators or independent events, AI services may reason over interconnected knowledge describing threat actors, vulnerabilities, attack techniques, organisational assets, defensive controls, regulatory obligations, and historical observations. Such contextual reasoning improves the ability to assess complex cybersecurity situations while supporting transparent analytical workflows that remain understandable to human analysts. The SKM is also designed

to support the emerging paradigm of agent-oriented artificial intelligence, since it supports the CyberAid LLM orchestration and agentic-analysis capabilities specified in D2.5, that require governed access to structured knowledge, contextual evidence and explicit semantic relationships in order to perform retrieval, analysis and reporting using a shared interpretation of cybersecurity concepts. Autonomous or semi-autonomous cybersecurity agents increasingly require structured knowledge environments through which they can discover relevant information, exchange contextual knowledge, coordinate analytical activities, and justify their recommendations. Because the SKM represents cybersecurity concepts through a consistent semantic framework, it provides a common knowledge layer that enables multiple intelligent agents to operate using a shared understanding of cybersecurity entities, relationships, governance policies, and operational context. This capability is particularly relevant for distributed cybersecurity environments where multiple analytical services cooperate to achieve common resilience objectives. Every retrieval result used as supporting evidence shall include a citable locator that resolves to the authoritative source element. For documentary knowledge, the locator shall identify the source document and the applicable section, page or fragment where available. For structured knowledge, it shall identify the canonical entity, relationship or graph node. The locator shall be returned together with the canonical SKM identifier, provenance reference, classification, confidence and temporal-validity metadata. This representation enables consuming components to verify that analytical claims remain grounded in referenced knowledge.

Another significant characteristic of the proposed representation concerns explainability. As AI becomes increasingly involved in cybersecurity decision support, it is essential that analytical recommendations remain transparent and justifiable. The explicit semantic structures maintained by the SKM allow analytical outcomes to be associated with the entities, relationships, observations, and evidence that contributed to a particular conclusion. Consequently, AI-generated recommendations may be traced back to the underlying knowledge from which they originated, improving analyst confidence while supporting governance, auditing, and regulatory compliance. The semantic representation additionally supports context-aware knowledge enrichment. As new cyber threat intelligence is integrated through the processes described in Chapter 4, existing semantic entities may be expanded with additional relationships, operational observations, financial context, regulatory information, or organisational knowledge. This continuous enrichment enables the SKM to evolve alongside the cybersecurity landscape while preserving semantic consistency and maintaining compatibility with AI-driven analytical services. The resulting knowledge environment therefore becomes progressively more valuable as intelligence accumulates over time. An important architectural characteristic of the proposed approach is its independence from individual AI technologies. The SKM is not designed specifically for a particular large language model, graph processing framework, reasoning engine, or autonomous agent platform. Instead, it provides a stable semantic representation that may support diverse AI methodologies as they evolve throughout the lifecycle of the CyberAid project. This technology-neutral design preserves long-term architectural flexibility while ensuring that future analytical capabilities can exploit the same underlying knowledge foundation without requiring fundamental modifications to the representation model itself.

Eventually, the semantic and AI-ready representation proposed by CyberAid transforms cybersecurity knowledge from a passive information repository into an active analytical resource capable of supporting intelligent retrieval, contextual reasoning, graph exploration, explainable AI, Retrieval-Augmented Generation, and future agent-oriented cybersecurity services. By combining formal semantic organisation with technology-independent knowledge structures, the SKM establishes the common representation layer required for advanced cybersecurity intelligence and decision support across the CyberAid architecture.

5.3 Knowledge Consumption Across CyberAid

The SKM constitutes the common semantic foundation upon which the analytical and operational capabilities of the CyberAid platform are built. Rather than functioning as an isolated knowledge repository, the SKM provides a shared representation of cybersecurity concepts, relationships, contextual information, governance attributes, and cyber threat intelligence that can be consistently interpreted by multiple platform components. This common knowledge layer enables diverse analytical services to operate upon

the same semantic understanding of the cybersecurity domain, thereby reducing fragmentation, improving interoperability, and supporting coherent decision-making across the CyberAid architecture.

Table 11 - Domain to consumer map

SKM semantic domain	Representative CyberAid consumer	Logical knowledge scope	Semantic graph scope
Threat and CTI	Email Threat Detection; Automated Log Analysis and Event Correlation	Threat-intelligence knowledge	Threats, actors, campaigns and indicators
Attack and adversarial behaviour	Email Threat Detection; AI-Driven Pentesting	Attack knowledge	Techniques, tools, malware and attack paths
Vulnerability and exposure	AI-Driven Pentesting; Code Analysis; Risk and Compliance Assessment	Vulnerability knowledge	Vulnerabilities, weaknesses, products and affected assets
Detection, observation and incident	Anomaly Detection; Automated Log Analysis and Event Correlation; Reporting	Operational evidence	Observations, findings, detections and incidents
Defensive knowledge	Pentesting Outcomes Analysis; Risk and Compliance Assessment	Defensive-control knowledge	Controls, mitigations and response measures
Financial and organisational context	Risk and Compliance Assessment	Organisational context	Assets, services, processes and dependencies
Regulatory and compliance	Risk and Compliance Assessment; Reporting	Regulatory knowledge	Requirements, obligations, controls and evidence
Metadata, provenance and governance	All authorised consumers	Governance metadata	Provenance, classification, lifecycle and ownership

The logical knowledge and graph scopes shown above are semantic groupings defined by D2.4, while D2.5 may bind them to runtime retrieval namespaces or Security Context fields. For example, the Vulnerability and exposure row assigns the same SKM semantic domain to the AI-Driven Pentesting, Code Analysis, and Risk and Compliance Assessment consumers. These consumers may retrieve vulnerabilities, weaknesses, products, and affected assets within an authorised logical knowledge scope. D2.4 defines that semantic scope, whereas D2.5 defines its binding to runtime namespaces, Security Context fields, and enforcement mechanisms.

A fundamental characteristic of the proposed architecture is the separation between knowledge production and knowledge consumption. Chapters 3 and 4 described how cybersecurity knowledge is modelled and how heterogeneous cyber threat intelligence is integrated into the SKM. Once this integrated knowledge has been established, individual CyberAid components consume it according to their respective functional responsibilities without requiring independent knowledge representations or duplicated integration processes. Consequently, the SKM acts as a central semantic reference that promotes consistency across the entire platform. One of the primary consumers of the SKM is cybersecurity monitoring and situational awareness. Monitoring components continuously interpret cybersecurity observations in relation to the semantic entities and relationships maintained within the knowledge model, allowing operational events to be analysed within their broader contextual environment. Instead of treating alerts or observations as isolated occurrences, monitoring capabilities may associate them with known threat actors, vulnerabilities, attack techniques, affected organisational assets, previous incidents, and contextual intelligence. This semantic interpretation enhances situational awareness while supporting more informed operational responses. The SKM also provides the knowledge foundation required for cyber-risk assessment. Risk-oriented components exploit the semantic relationships connecting technical vulnerabilities, threat intelligence, organisational assets, financial processes, defensive controls, and regulatory obligations to evaluate cybersecurity exposure from multiple perspectives. By consuming a common semantic representation, risk analysis can consider not only individual technical weaknesses but also their potential operational and business implications within the financial sector. This integrated perspective enables more comprehensive and context-aware risk evaluation than would be achievable through isolated technical assessments.

Another important category of knowledge consumers comprises decision-support services. Strategic and operational cybersecurity decisions frequently require the integration of technical intelligence, organisational priorities, regulatory considerations, financial impacts, and historical observations. The SKM enables decision-support components to access these heterogeneous information domains through a unified semantic structure, allowing analytical results to be generated from interconnected knowledge rather than disconnected datasets. This capability supports more transparent prioritisation, facilitates evidence-based recommendations, and strengthens the consistency of cybersecurity decision-making processes. The semantic representation further supports AI-assisted analytical services, which consume structured knowledge to perform contextual retrieval, intelligent reasoning, knowledge exploration, and advanced analytical assistance. Because cybersecurity concepts, relationships, provenance information, and governance attributes are explicitly represented within the SKM, AI services can interpret cybersecurity situations using structured contextual information rather than relying exclusively on unstructured textual content. This shared semantic foundation improves analytical consistency while enabling AI-generated outputs to remain grounded in authoritative cybersecurity knowledge. Within the CyberAid architecture, this also includes the Email Threat Detection capability (EMTD), which may consume SKM knowledge to obtain relevant threat context, semantic relationships, indicators, historical evidence, and other contextual information required to support its analytical processes.

The SKM additionally supports graph-based analytical capabilities, where the primary objective is to explore the relationships connecting cybersecurity entities across multiple semantic domains. Graph-oriented analyses consume the interconnected knowledge representation to investigate attack paths, identify infrastructure dependencies, analyse behavioural patterns, discover indirect relationships, and evaluate multi-stage attack scenarios. Since these analyses rely upon explicit semantic connections maintained within the SKM, they benefit directly from the contextual integration achieved throughout the preceding chapters. Another important consumer of the SKM is the regulatory and compliance assessment capability. Cybersecurity activities within the financial sector operate under an evolving regulatory landscape that includes governance requirements, resilience obligations, reporting responsibilities, and security control expectations. By integrating regulatory concepts alongside technical cybersecurity knowledge, the SKM enables compliance-oriented components to interpret operational observations within their applicable governance context. This semantic integration facilitates more consistent compliance assessment while supporting explainable regulatory decision support. Regulatory reporting capabilities

may additionally consume the SKM to identify applicable reporting obligations, associated supporting evidence, competent reporting contexts, and relevant notification or reporting timelines linked to cybersecurity incidents and regulatory requirements. In this way, regulatory deadlines and reporting obligations can be represented as contextual knowledge supporting consistent and explainable compliance-oriented decision support. The representation also provides a common knowledge environment for future autonomous and collaborative cybersecurity services. As intelligent analytical components increasingly operate through cooperative workflows, access to a shared semantic understanding becomes essential for maintaining consistency across distributed decision-making processes. The SKM provides this common contextual reference, allowing multiple services to exchange information, interpret cybersecurity observations, and coordinate analytical activities using the same conceptual vocabulary and relationship structures. This architectural characteristic strengthens interoperability while supporting the long-term evolution of CyberAid towards increasingly intelligent and collaborative cybersecurity capabilities. A significant advantage of the proposed consumption model is that all platform capabilities operate upon the same continuously evolving knowledge foundation. As additional cyber threat intelligence is integrated through the mechanisms described in Chapter 4, newly acquired knowledge becomes immediately available to every authorised knowledge consumer without requiring independent integration or transformation processes. This shared knowledge paradigm reduces redundancy, improves analytical consistency, and ensures that all components benefit from the same contextual understanding of the cybersecurity environment while respecting applicable governance and access-control policies.

Although different CyberAid components consume the SKM for different analytical objectives, they all rely upon the same semantic principles, conceptual entities, contextual relationships, provenance information, and governance attributes defined throughout this deliverable. This common semantic foundation enables independent services to produce complementary analytical outcomes while maintaining architectural consistency across the platform. Consequently, the SKM functions not simply as a repository of cybersecurity information, but as the semantic integration layer that connects the diverse analytical capabilities of CyberAid into a coherent and interoperable cybersecurity ecosystem.

6. CONCLUSIONS AND NEXT STEPS

This deliverable has presented the conceptual specification of the CyberAid SKM and the Cyber Threat Intelligence (CTI) Integration Framework, establishing the semantic foundations required to support knowledge-driven cybersecurity capabilities throughout the CyberAid platform. The proposed specification addresses the challenge of integrating heterogeneous cybersecurity information originating from multiple technical, organisational, financial, and regulatory domains into a unified semantic architecture that promotes interoperability, contextual understanding, and intelligent analytical processing. The SKM introduced in this deliverable provides a structured conceptual representation of cybersecurity knowledge through well-defined semantic domains, relationships, metadata structures, provenance information, governance attributes, and lifecycle principles. Rather than focusing on individual technologies or implementation mechanisms, the specification defines a technology-independent semantic architecture capable of representing complex cybersecurity ecosystems while remaining sufficiently flexible to accommodate future developments, evolving threat landscapes, and emerging analytical requirements. Complementing the SKM, the CTI Integration Framework specifies the logical processes through which externally acquired cyber threat intelligence is transformed into contextualised organisational knowledge. The framework defines a complete integration workflow encompassing intelligence acquisition, validation, normalisation, semantic correlation, contextual enrichment, semantic mapping, governance, and quality assurance. Together, these processes establish a consistent mechanism for integrating heterogeneous cyber threat intelligence while preserving semantic meaning, provenance, contextual information, and interoperability with recognised intelligence exchange standards.

The deliverable has further demonstrated that the resulting knowledge architecture is suitable for multiple complementary computational representations, including semantic models, ontology-oriented structures, and graph-oriented knowledge representations. By separating conceptual modelling from implementation

technologies, the proposed architecture enables different computational paradigms to coexist while preserving the consistency of the underlying semantic model. This architectural flexibility strengthens interoperability and provides a robust foundation for future technological evolution without requiring modifications to the conceptual specification itself. An important contribution of the proposed architecture is its readiness to support advanced analytical capabilities increasingly required within modern cybersecurity platforms. The semantic organisation of cybersecurity knowledge, together with the explicit representation of contextual relationships, enables graph analytics, semantic retrieval, AI-assisted reasoning, Retrieval-Augmented Generation (RAG), explainable decision support, and future agent-oriented analytical services to operate upon a common knowledge foundation. Consequently, the SKM serves not only as a repository of cybersecurity knowledge but also as the semantic integration layer supporting coherent analytical workflows across the CyberAId ecosystem. The specification also reinforces the importance of governance, provenance, and information quality as fundamental characteristics of cybersecurity knowledge management. By incorporating lifecycle management, confidence assessment, source traceability, dissemination policies, and governance attributes into the conceptual architecture, the proposed approach ensures that integrated knowledge remains reliable, explainable, auditable, and suitable for operational decision-making within the highly regulated environment of the financial sector.

Overall, the work presented in this deliverable establishes a coherent architectural framework connecting cybersecurity knowledge representation, cyber threat intelligence integration, semantic interoperability, and intelligent knowledge consumption. These complementary elements collectively provide the conceptual basis upon which subsequent CyberAId technical developments can be designed, implemented, and evaluated while maintaining consistency across the different work packages contributing to the platform. The activities described in this deliverable represent the architectural specification phase of the CyberAId knowledge management capabilities. Building upon these foundations, the subsequent phases of the project will progressively realise the proposed concepts through prototype development, technical integration, validation activities, and cross-work-package collaboration. Future work will focus on refining the SKM, extending semantic mappings as additional intelligence sources become available, strengthening interoperability with other CyberAId components, and validating the proposed architecture within representative financial-sector cybersecurity scenarios. These activities will be reported in the subsequent technical deliverables of the project, ensuring continuity between the conceptual design established in D2.4 and its progressive implementation throughout the remaining project lifecycle.

In conclusion, the SKM and the CTI Integration Framework provide the semantic backbone of the CyberAId architecture. By establishing a technology-independent, interoperable, and context-aware approach to cybersecurity knowledge management, this deliverable creates the foundations necessary for intelligent cybersecurity services capable of supporting advanced threat intelligence integration, explainable analytics, collaborative decision support, and resilient cyber defence across the financial sector.

REFERENCES

- [1] CyberAId D2.1 – Financial Sector Cybersecurity Requirements Analysis
- [2] CyberAId D2.2 – Reference Architecture and Blueprints – First Version
- [3] CyberAId D2.5 – LLM-based Cybersecurity Orchestration Design
- [4] [OASIS Open](#), “STIX Version 2.1,” OASIS Standard, 10 June 2021
- [5] [OASIS Open](#), “TAXII Version 2.1,” OASIS Standard, 10 June 2021
- [6] [MITRE](#), “MITRE ATT&CK Knowledge Base.”
- [7] [CVE Program](#), “Common Vulnerabilities and Exposures.”
- [8] [National Institute of Standards and Technology](#), “National Vulnerability Database.”
- [9] [FIRST](#), “Common Vulnerability Scoring System Version 4.0: Specification Document.”
- [10] [MITRE](#), “Common Weakness Enumeration.”
- [11] [MITRE](#), “Common Attack Pattern Enumeration and Classification.”
- [12] [FIRST](#), “Traffic Light Protocol: Version 2.0”
- [13] [OpenCTI](#), “OpenCTI Documentation.”
- [14] [FS-ISAC](#), “Operating Rules”.
- [15] C. Johnson, M. Badger, D. Waltermire, J. Snyder and C. Skorupka, “[Guide to Cyber Threat Information Sharing](#),” NIST Special Publication 800-150, October 2016