



Introduction

Handle Complexity

Ontology as the Foundation
of Multi-Compliance

AI and Automation

Collaboration and
Consensus

Roadmap

Methodology and
Design Principles

Conclusion: Ontology as an
Enabler of Resilient
Multi-Compliance Capability

Ontology enhanced multi-compliance capability

Executive Summary



Cybersecurity is becoming a business challenge as much as a technical one. Organizations are confronted with a growing number of regulations, standards, customer requirements, and supply-chain obligations. At the same time, cyber risks, digital dependencies, and management accountability **continue to increase**. The resulting complexity makes it increasingly difficult to understand how controls, risks, technologies, services, and compliance requirements relate to one another. This paper proposes ontology-based enhancement of the ZOIS Security Pillar as a foundation for machine-readable, multi-compliance cybersecurity governance.

The ZOIS Security Pillar, based on the ESARIS taxonomy, already provides a strong basis for **organizing security knowledge**. However, taxonomies are primarily designed for human navigation and classification. They are less effective at representing the relationships between concepts and are difficult to integrate into modern automation, compliance, and AI-driven processes.

This paper proposes the development of ontology as the next evolution of the **ESARIS-based security framework**. An ontology extends a taxonomy by introducing machine-readable definitions and explicit relationships between concepts such as controls, assets, risks, services, suppliers, business processes, and regulatory requirements. The result is a connected knowledge model that can be used by both humans and machines.

The primary value of such a model lies in its ability to support multi-compliance. Rather than interpreting each regulation or standard independently, organizations can describe their architecture and security environment once, then map that reality to multiple compliance frameworks simultaneously. This creates traceability, improves consistency, and reduces duplication of effort.



An ontology also creates the foundation for advanced capabilities such as automated compliance analysis, decision support, AI-assisted governance, structured reasoning, and future Digital Twin applications. As cybersecurity knowledge becomes machine-readable and semantically connected, organizations gain the ability to understand dependencies, assess the impact of change, and navigate increasingly complex governance environments more effectively.

The paper further argues that the success of such an initiative depends on collaboration. An ontology is ultimately a formal representation of shared understanding. Its long-term value is derived not from technology alone, but from industry consensus regarding concepts, terminology, responsibilities, and relationships. As an industry association focused on secure and resilient supply chains, ZOIS is uniquely positioned to facilitate this collaboration.

A practical roadmap is proposed, beginning with machine-readable representations of standards, continuing with relationship modeling and user-facing tools, and culminating in operational adoption by participating organizations. Individual organizations can extend the shared ontology to reflect their own environments while maintaining interoperability with partners, customers, and suppliers.

The central message of this paper is straightforward: ontology technology is not an end in itself. It is an enabling capability that can help ZOIS transform existing security knowledge into a machine-readable, interoperable, and collaborative foundation for multi-compliance, automation, and supply-chain resilience. By doing so, it can strengthen the ability of organizations to manage cybersecurity complexity and work together more effectively in an increasingly interconnected world.

1. Introduction



Growing complexity of cybersecurity

Cybersecurity has evolved far beyond its traditional role as a technical function supporting information technology operations. Today, it has become a business-critical discipline that directly influences organisational resilience, regulatory compliance, operational continuity, and corporate reputation. Several converging trends are driving a sustained increase in complexity across the cybersecurity domain.

First, the regulatory environment has fundamentally changed. Within the European Union, cybersecurity is increasingly treated as a governance and business responsibility rather than merely a technical concern. The NIS2 Directive explicitly introduces management accountability for cybersecurity risk management, requiring executives and boards to approve, oversee, and be accountable for cybersecurity measures. Cybersecurity has therefore moved into the boardroom and become an element of corporate governance comparable to financial, legal, or operational risk management¹.

Second, the financial consequences of cybersecurity failures continue to increase. According to IBM's Cost of a Data Breach Report 2025², the global average cost of a data breach reached approximately USD 4.44 million. Ransomware-related incidents exceeded USD 5 million on average, while organisations frequently experienced prolonged operational disruption and recovery periods following major incidents. Beyond direct remediation costs, organisations increasingly face legal liability, regulatory penalties, reputational damage, and customer attrition.

¹ <https://digital-strategy.ec.europa.eu/en/policies/nis2-directive>

² <https://www.ibm.com/reports/data-breach>



Third, organisations operate within increasingly interconnected digital ecosystems. Cloud adoption, outsourcing, managed services, software supply chains, and platform-based business models have significantly expanded the attack surface. Security decisions can no longer be evaluated solely within organisational boundaries; they must account for dependencies on suppliers, partners, service providers, and customers. Consequently, cybersecurity has become a systems problem involving multiple stakeholders, overlapping responsibilities, and complex dependency networks.

Finally, the volume of information that security leaders must process continues to grow. Security architecture must simultaneously accommodate technical threats, compliance obligations, risk-management processes, business objectives, and evolving operational requirements. The challenge is no longer simply implementing controls; it is understanding the relationships, dependencies, and consequences that exist between those controls across a highly dynamic environment.

Explosion of standards, frameworks, and regulations

The increasing complexity of cybersecurity is further amplified by the rapid growth of standards, frameworks, regulations, and industry-specific requirements. Organisations no longer operate under a single compliance regime. Instead, they must simultaneously satisfy multiple regulatory and contractual obligations that frequently overlap, partially duplicate one another, or use different terminology to describe similar security objectives.

This challenge is particularly relevant within the European context. Mario Draghi's 2024 report on the future of European competitiveness³ identifies regulatory fragmentation and increasing compliance burdens as factors affecting the competitiveness of European industry. The report highlights the need to simplify regulatory interactions and improve the effectiveness of governance frameworks supporting the digital and economic transformation of Europe.

³ https://commission.europa.eu/topics/competitiveness/draghi-report_en



In the cybersecurity domain alone, organisations may need to address obligations arising from NIS2, the Cyber Resilience Act, DORA, GDPR, sector-specific requirements, national regulations, customer security requirements, contractual obligations, and internationally recognised standards such as ISO/IEC 27001, NIST Cybersecurity Framework, CIS Controls, and industry-specific frameworks. Each introduces additional requirements, evidence expectations, reporting obligations, and governance structures⁴.

The challenge becomes even more pronounced in supply-chain environments. A core objective of the Zero Outage Industry Standard (ZOIS) is to strengthen security across interconnected supply chains. In such environments, regulatory obligations do not remain isolated within a single organisation. Instead, requirements propagate through supplier ecosystems, where organisations inherit security expectations from customers, industry bodies, regulators, and strategic partners. NIS2 itself explicitly introduces supply-chain security obligations and risk-management requirements extending beyond organisational boundaries.

As a result, compliance requirements become additive rather than substitutive. A supplier serving multiple industries or jurisdictions may need to demonstrate compliance with numerous overlapping frameworks simultaneously. The resulting governance landscape resembles a complex network of interconnected obligations rather than a collection of independent standards. Understanding how these requirements align, overlap, reinforce, or conflict with one another increasingly exceeds the capacity of purely manual approaches.



Need for a better navigation model

The combined effect of increasing business importance and growing regulatory complexity creates a fundamental scalability challenge. Human-readable documents remain essential for communication and governance, but they are increasingly insufficient as the primary mechanism for managing cybersecurity knowledge. Modern security programmes require machine-assisted analysis to maintain visibility across thousands of controls, requirements, assets, risks, and compliance relationships.

Recognising this challenge, the National Institute of Standards and Technology (NIST) launched the Open Security Controls Assessment Language (OSCAL) initiative⁵. OSCAL provides machine-readable representations of security controls, implementations, assessments, and compliance artefacts using standardised data formats. The objective is to transform traditional document-centric compliance processes into data-centric processes that support automation, interoperability, and continuous assessment. NIST explicitly positions OSCAL as a mechanism to automate security and compliance activities and to convert security-control information into machine-readable formats⁶.

The move toward machine-readable controls represents a critical first step, but it does not by itself solve the problem of understanding complex governance systems. Digital representations of controls primarily describe individual elements. To generate meaningful insights, recommendations, predictions, or automated reasoning, systems must also understand the relationships between those elements.

⁵ <https://pages.nist.gov/OSCAL/>

⁶ <https://csrc.nist.gov/Projects/Open-Security-Controls-Assessment-Language>



This requirement points toward a more advanced model: the representation of cybersecurity knowledge as a connected network of concepts, dependencies, responsibilities, controls, risks, assets, and regulatory requirements. In other words, cybersecurity frameworks must evolve from static taxonomies toward richer semantic structures capable of describing not only what exists, but also how individual elements interact with one another.

Such a model functions as a navigational layer for increasingly complex governance environments. It enables organisations to trace compliance obligations across multiple frameworks, understand dependencies between controls, identify the consequences of change, and support machine-assisted analysis. As cybersecurity complexity continues to grow, this ability to explicitly represent relationships may become as important as the controls themselves. The next chapter explores how ontologies can provide such a navigation model and serve as the foundation for enhanced multi-compliance capability.

2. Handle Complexity



2.1

Why traditional taxonomies do not scale

The ZOIS Security Pillar is built upon the ESARIS (Enterprise Security Architecture for Reliable ICT Services) framework, which provides a structured taxonomy for organising security-related standards, controls, responsibilities, and activities. Within the broader ZOIS architecture, the Security workstream is explicitly based on the ESARIS Security Taxonomy, positioning it as the primary organisational structure for security knowledge and guidance.

This choice provides significant advantages. Taxonomies create order. They establish a common vocabulary, organise knowledge into understandable categories, support communication between stakeholders, and help organisations assign responsibilities and orchestrate security activities. As the authors of the Taxonomy for Systematic IT Service Security Along the Entire Supply Chain⁷ observe, the primary purpose of a taxonomy is to provide an overview, unify terminology, guide collaboration, and help organise security activities across organisations and supply chains.

These characteristics have made ESARIS highly successful as a governance instrument. Security practitioners can locate standards, understand ownership, identify relevant activities, and navigate the security domain through a coherent structure. The taxonomy acts as a map of the security landscape and thereby reduces complexity. As von Faber notes, complexity is the opponent of effective cybersecurity, and taxonomies help make complexity manageable by providing structure and supporting communication.

7

<https://zero-outage.com/news/taxonomy-for-systematic-it-service-security-along-the-entire-supply-chain/>



However, the strengths of a taxonomy also reveal its limitations.

A taxonomy primarily answers the question "Where does this belong?". It is essentially a classification system. While it organises concepts into categories and hierarchies, it says relatively little about how those concepts interact. Controls, risks, technologies, services, suppliers, regulations, responsibilities, and business objectives can all be classified, but their dependencies remain largely implicit.

Furthermore, most taxonomies are consumed as documents rather than as digital knowledge structures. They are designed for human reading, not machine processing. Relationships between elements often exist only in prose, diagrams, or expert interpretation. Consequently, integrating taxonomies with management systems, governance tools, AI-based analysis, regulatory mapping, or automated compliance platforms requires considerable manual effort.

This limitation becomes increasingly problematic as the volume of requirements, standards, and compliance obligations grows. The challenge is no longer simply finding information within the taxonomy. Rather, it is understanding how hundreds or thousands of elements influence one another across a dynamic ecosystem of security, architecture, operations, suppliers, and regulations.

The question therefore becomes: **how do we preserve the clarity of a taxonomy while enabling navigation through a far more complex landscape?**



2.2

The London Underground lesson: why maps matter

A useful analogy can be found outside cybersecurity.

Harry Beck's famous solution (Figure 1 - London Underground map) was revolutionary because he stopped attempting to represent physical reality and instead focused on representing relationships. Stations were simplified, distances were distorted, and lines were straightened. The resulting map was not geographically accurate, but it dramatically improved people's ability to reason about the transportation system.

The railway itself did not change. The mental model did.

Beck's innovation reduced cognitive load by creating a stable schema through which users could understand complexity, adapt to change, and make decisions more effectively. The network continued to grow, but passengers no longer needed to continuously relearn how it worked.

Modern cybersecurity faces a remarkably similar challenge.



Figure 1 - London Underground map



The security landscape consists of controls, requirements, technologies, responsibilities, assets, suppliers, regulations, risks, and processes connected through countless dependencies. Traditional document structures often describe these elements accurately but provide limited support for navigating their interactions. The result is a growing cognitive burden placed upon security professionals, architects, auditors, and business decision-makers.

This is where ontologies become interesting.

An ontology is not simply a visualisation tool and should not be confused with a knowledge graph. However, ontologies can provide the semantic foundation from which graph-based representations, relationship maps, decision-support models, and navigation tools can be derived. In this sense, an ontology performs a role similar to Beck's Underground map: it provides a simplified and structured representation of relationships that allows humans and machines to reason more effectively about a complex system.



2.3

From taxonomy to ontology

The distinction between a taxonomy and an ontology is subtle but important.

A taxonomy classifies concepts into hierarchical categories. It primarily expresses parent-child relationships and provides an organisational structure for information.

An ontology extends this idea considerably. In addition to classification, it formally defines concepts, attributes, constraints, and relationships between concepts. Rather than merely specifying where an element belongs, an ontology describes how elements interact. For example, it can formally express that a control mitigates a risk, a supplier provides a service, a regulation mandates a requirement, or a technology supports a business capability.

In practical terms, transforming a taxonomy into an ontology requires several additions:



Formal definitions of concepts



Explicit relationships between concepts



Consistent semantics and terminology



Machine-readable representation formats



Rules and constraints that allow reasoning over the resulting model



The result is not simply a better-organised document, but a structured knowledge system that can be queried, analysed, reused, and integrated into digital workflows.

Historically, however, developing ontologies has been a difficult undertaking. Ontology engineering involves significant effort in requirements analysis, concept modelling, relationship discovery, semantic harmonisation, validation, and governance. Building consensus among experts is often more difficult than implementing the technology itself. The process can require years of work and continuous maintenance.

Fortunately, recent technological developments are changing this picture. Artificial intelligence, natural language processing, and machine-assisted knowledge extraction can now help identify concepts, suggest relationships, detect inconsistencies, and accelerate ontology development activities. While human expertise remains indispensable, the cost and effort required to construct and maintain ontologies are decreasing rapidly.

We will return to this topic in Chapter 4 when discussing the relationship between ontologies, AI, and automation.



2.4

Ontologies as security mindmaps

Mindmap can be used as an analogy for understanding ontologies. This comparison should not be taken literally. An ontology is far more rigorous than a mindmap and includes formal semantics, machine-readability, and logical structure that go well beyond simple visual diagrams.

Nevertheless, the analogy is useful. A mindmap helps people understand complex domains by showing concepts and connections in a single navigable structure. Similarly, an ontology provides a coherent representation of a knowledge domain that makes relationships visible and allows both people and technology to navigate complexity.




Figure 2 - Ontology provides the framework to tackle the challenge ahead of us



This capability becomes particularly relevant when considering another recent ZOIS contribution: Taxonomy for Systematic IT Service Security Along the Entire Supply Chain. (Figure 2 - Ontology provides the framework to tackle the challenge ahead of us) The paper proposes a significantly evolved taxonomy structure that organises security activities around practices, technologies, supply-chain interactions, lifecycle stages, and operational responsibilities. Rather than focusing solely on technical controls, it provides an enterprise-oriented view of security management and service delivery across interconnected organisations.

The transition from the original ESARIS taxonomy towards such a richer model presents a significant challenge. New categories must be introduced, concepts must be harmonised, and relationships between previously isolated areas must be made explicit.

An ontology provides the mechanism to facilitate this evolution.

Instead of manually rebuilding taxonomies whenever new organisational structures emerge, existing concepts and controls can be linked through a semantic layer. New viewpoints, classifications, and navigational structures can then be generated from the same underlying knowledge model. The ontology becomes a stable foundation upon which multiple taxonomies, visualisations, compliance mappings, and reporting structures can coexist.



Beyond supporting taxonomy evolution, the business benefits extend further:

- Improved understanding of control dependencies
- Better impact analysis when requirements change
- Traceability across standards and regulations
- Enhanced collaboration across organisational boundaries
- Greater support for automation and AI-driven analysis
- More effective management of complex supply-chain ecosystems
- Faster onboarding of new stakeholders
- Stronger multi-standard and multi-regulatory alignment

Taken together, these capabilities suggest that the value of an ontology extends well beyond knowledge organisation. It becomes a foundational mechanism for managing complexity itself.

This observation leads directly to the next question: if ontologies can connect controls, requirements, responsibilities, technologies, and regulations into a unified semantic model, can they also provide the foundation for managing compliance across multiple standards simultaneously? The following chapter explores precisely this possibility and examines how ontology-based approaches can enable enhanced multi-compliance capabilities.



3. Ontology as the Foundation of Multi-Compliance

The previous chapter argued that taxonomies are effective for organising information but increasingly struggle to support decision-making in environments characterised by growing complexity, overlapping requirements, and continuous change. Ontologies address this limitation by moving beyond classification and introducing explicit, machine-readable relationships between concepts. This capability is not merely a technical improvement; it creates the foundation for a new approach to managing compliance itself.

Traditionally, organisations implement controls to satisfy individual standards and regulations. As the number of applicable frameworks increases, controls are mapped repeatedly against different requirements, often resulting in duplicated effort and inconsistent interpretations. An ontology enables a different perspective. Instead of treating compliance frameworks as the primary objects of analysis, it places the organisation's architecture, assets, processes, risks, and controls at the centre, allowing multiple compliance views to be generated from a common underlying model.

3.1

Mapping controls, risks, assets and requirements

The key advantage of an ontology lies in its ability to represent security knowledge in a machine-readable form. Unlike traditional documents, ontology concepts are not merely described; they are connected through formally defined relationships that can be processed by software systems.

This enables a direct connection between security governance and the actual environment being governed.



Every organisation already possesses an architectural reality consisting of applications, infrastructure components, data repositories, identities, suppliers, business services, security technologies, and operational processes. Compliance requirements are ultimately applied to these elements. Whether a requirement originates from ISO 27001, NIS2, DORA, customer contracts, or industry-specific standards, it is the same architectural landscape that must be secured.

An ontology makes it possible to model these architectural elements together with their associated controls, risks, obligations, and ownership structures. Assets can be linked to business services, services to processes, processes to risks, risks to controls, and controls to regulatory requirements. Once these relationships are represented explicitly, the organisation gains something that traditional control catalogues cannot provide: a semantic model of its security architecture.

The practical consequence is significant. Instead of maintaining separate interpretations of multiple standards, the organisation can describe its environment once and derive multiple compliance perspectives from that description. A backup service remains a backup service regardless of which framework references it. An identity management capability remains the same capability whether viewed through the lens of ISO 27001, NIS2, SOC2, or a customer-specific security assessment.

In this sense, the ontology functions as an intermediary layer between operational reality and compliance obligations. The description of the IT environment becomes translatable into different compliance control sets because the controls themselves are connected to a common conceptual model rather than being treated as isolated requirements.

This shift from document-centric compliance to architecture-centric compliance forms the basis of a scalable multi-compliance capability.



3.2

Navigating overlapping standards

If an ontology provides the common language, mappings provide the translation mechanism.

Control mapping is not a new challenge. Organisations have long attempted to create crosswalks between standards, frameworks, and regulations. However, these efforts are frequently limited to simple references showing that one control is somehow related to another. Such mappings often lack context regarding the nature of the relationship or its practical significance.

Recognising this challenge, NIST published NIST IR 8477: Mapping Relationships Between Documentary Standards, Regulations, Frameworks, and Guidelines⁸. The report introduces a structured methodology for mapping cybersecurity concepts across different sources and emphasises that concepts such as controls, requirements, technologies, functions, processes, and outcomes should not be viewed as isolated elements but as components of a broader concept system.

Particularly relevant is NIST's observation that concepts do not exist independently and that useful mappings must explicitly describe the relationships between them. The methodology distinguishes several types of relationships, including equivalent concepts, supportive relationships, parent-child structures, overlapping concepts, and prerequisite relationships. Rather than creating simple cross-references, the approach seeks to capture why two concepts are connected and how one contributes to another.

⁸ <https://csrc.nist.gov/pubs/ir/8477/final>



This philosophy aligns closely with the ontology-based approach proposed in this paper.

Whereas traditional mappings are often maintained manually as spreadsheets, an ontology can incorporate these relationships directly into the underlying knowledge model. As additional standards or regulations emerge, they can be integrated through relationships to existing concepts rather than through isolated compliance projects. This becomes particularly valuable in supply-chain environments where multiple organisations, customers, regulators, and industry frameworks simultaneously influence security requirements.

The result is not the elimination of compliance complexity, but its transformation into a navigable structure.

3.3

Decision support

The most powerful consequence of an ontology may ultimately be neither compliance management nor control mapping, but decision support.

A conventional taxonomy provides a flat classification structure. It tells users where a concept belongs. An ontology adds relationships and dependencies, allowing the same information to be navigated as a decision structure.

An analogy can be drawn to the difference between a list of destinations and a route planner. Both contain information about locations, but only the latter can help determine the consequences of choosing one path over another.



Security architecture decisions exhibit similar characteristics.

Consider a seemingly simple architectural choice: introducing a cloud-based identity provider. Such a decision affects access management, logging requirements, supplier assurance activities, data residency considerations, business continuity plans, audit evidence collection, and potentially multiple compliance obligations. Traditional documentation can describe these topics separately, but understanding the downstream consequences often relies on expert interpretation.

An ontology offers the possibility of making these dependencies explicit.

If controls, assets, risks, services, suppliers, technologies, and compliance requirements are connected through formally described relationships, then architectural decisions can be analysed in the context of those relationships. Changes to one part of the model can be traced to the associated controls, affected risks, impacted compliance obligations, and potentially even to required governance actions.

This transforms security knowledge from a static repository into a navigable decision framework.

Importantly, the objective is not to automate architectural judgement or replace human expertise. Rather, it is to augment decision-making by making relevant dependencies visible and accessible. The ontology effectively acts as a security "mindmap" at organisational scale: not merely showing what exists, but revealing how elements influence one another and what consequences may emerge when changes occur.

This capability becomes even more valuable when combined with modern automation and AI technologies. Once relationships are represented in a machine-readable form, they become available for advanced analysis, reasoning, recommendations, and continuous compliance monitoring. The next chapter explores this connection and examines how ontologies can serve as a foundation for AI-assisted cybersecurity governance and automation.

4. AI and Automation



The previous chapter argued that ontologies can serve as the foundation for multi-compliance by connecting architectural reality with controls, risks, and regulatory requirements. However, the value of such a model is not limited to improving compliance management. Once security knowledge is represented in a structured, machine-readable form, it becomes available for a much broader set of automation and analytical capabilities.

This is where the discussion naturally intersects with artificial intelligence. AI systems can only reason effectively about information that they can access, interpret, and relate to other information. The more structured and explicit the underlying knowledge model becomes, the more effectively automation and AI can support security operations. Ontologies therefore should not be viewed as competitors to AI. Rather, they provide the semantic foundation that allows AI and automation technologies to operate with greater transparency, consistency, and traceability.

4.1

Machine-readable security knowledge

Most security standards are still delivered as documents intended for human consumption. While this format remains essential for communication, governance, and auditability, it creates limitations when organizations attempt to automate compliance activities.



The cybersecurity industry has already recognized this challenge. Governance, Risk, and Compliance (GRC) platforms such as ServiceNow GRC, RSA Archer, MetricStream, IBM OpenPages, and similar solutions increasingly rely on structured representations of controls, requirements, risks, assets, and evidence rather than purely document-based approaches. Likewise, modern cloud-security platforms, compliance monitoring tools, and security configuration assessment systems increasingly consume machine-readable control catalogs to automate assessments and reporting.

A notable example is NIST's Open Security Controls Assessment Language (OSCAL), which was specifically created to transform traditionally document-centric control frameworks into machine-readable formats. OSCAL allows controls, implementation statements, assessment plans, findings, and compliance artifacts to be expressed as structured data rather than static text. This enables automation, interoperability between tools, and continuous assessment approaches that would be difficult to achieve using documents alone.

For an organization such as ZOIS, whose purpose is to provide an industry standard, this development is particularly relevant. Future standards will increasingly be expected to support not only human readers but also software systems. A standard that exists exclusively as a PDF document risks becoming isolated from the growing ecosystem of automated governance, compliance, and security-management tools.

Machine-readable delivery should therefore be viewed as an important strategic objective rather than a technical enhancement. If standards are expected to participate in automated assessments, continuous compliance programs, digital audits, and AI-assisted governance processes, they must be represented in formats that machines can interpret and process effectively.



4.2

Linked data and connected security models

Machine readability alone, however, is not sufficient.

A collection of machine-readable control descriptions is valuable, but most organizational questions require understanding relationships rather than isolated facts. Security professionals rarely ask questions about a single control in isolation.

Instead, they seek to understand dependencies:



Which assets are affected by a particular vulnerability?



Which business services depend on those assets?



Which regulations reference the affected controls?



Which suppliers participate in the impacted service chain?



Which remediation activities must be completed first?

These questions require connected information rather than standalone records.



Ontologies address this requirement by enabling linked data models. Assets, controls, services, risks, suppliers, technologies, architectures, and regulatory obligations become connected objects within a common knowledge structure. The resulting model can be explored from multiple viewpoints and reused across operational, governance, and strategic activities.

One particularly interesting future application is the emerging concept of the Digital Twin. While Digital Twin technology originated in engineering and industrial systems, its underlying principle—a dynamic digital representation of a real-world environment—has increasing relevance for cybersecurity. A security ontology could provide part of the semantic foundation for a digital representation of an organization's security architecture, connecting technical components, governance structures, operational processes, and compliance requirements within a single navigable model.

Such a capability could support impact analysis, change management, compliance assessments, architecture planning, and operational decision-making. More importantly, it would allow organizations to simulate and evaluate changes before they are implemented in production environments.

Although this vision is still evolving, it demonstrates the broader potential of ontology-driven models beyond traditional compliance activities.



4.3

Structured reasoning

The true power of ontologies emerges when relationships are combined with reasoning.

Traditional databases store information. Ontologies can additionally support inference.

Reasoning engines evaluate the defined relationships between concepts and derive conclusions that have not been explicitly recorded. This capability is widely used in domains such as healthcare, engineering, logistics, and scientific research, where understanding relationships is often more valuable than storing individual facts.

In cybersecurity, a simple example might involve access management.

If an ontology defines that:



A privileged account has access to a particular system



The system processes regulated data



The regulation requires enhanced monitoring




Monitoring controls are not implemented

then a reasoning engine may infer the existence of an unresolved compliance gap even if that conclusion was never explicitly stated.



Similarly, reasoning can help identify:



From a business perspective, this transforms security governance from passive documentation into active decision support.

Another useful way to understand this shift is through the progression from a taxonomy to a decision tree. A taxonomy tells us where concepts belong. An ontology can help determine what follows from a particular choice.

Consider an organization evaluating two architectural alternatives: implementing a security capability internally or consuming it as a managed service. Both options can satisfy specific security objectives. However, each introduces different supplier dependencies, compliance obligations, operational responsibilities, contractual requirements, and risk profiles.

When such relationships are explicitly modelled, reasoning systems can assist architects, auditors, and security managers in understanding the consequences of their decisions before implementation occurs.

The result is not automated decision-making, but informed decision-making.

Organizations gain the ability to navigate increasingly complex governance environments while maintaining visibility into why particular conclusions were reached.



4.4

The role of AI in ontology development and operations

The relationship between AI and ontologies is complementary rather than competitive.

From an operational perspective, AI systems generally perform better when they operate within a structured context. Large Language Models are highly capable at synthesizing information, generating explanations, and assisting users. However, governance and compliance activities often require consistency, traceability, and adherence to explicit rules.

Ontologies can provide this structure.

Instead of asking an AI system to interpret hundreds of pages of standards independently, organizations can provide a semantic model containing formally defined concepts, relationships, responsibilities, constraints, and mappings. The AI can then use this structured foundation to generate recommendations, answer questions, explain dependencies, and support decision-making while remaining aligned with the organization's governance model.

In this sense, ontologies help make AI outputs more explainable and easier to validate.

The relationship also works in the opposite direction.

Historically, ontology development has required extensive manual effort. Subject-matter experts needed to identify concepts, define terminology, discover relationships, reconcile inconsistencies, and maintain the resulting knowledge model. This process often became a significant obstacle to adoption.

Recent advances in AI offer an opportunity to accelerate this work.

⁸ <https://csrc.nist.gov/pubs/ir/8477/final>



AI systems can already extract concepts from documents, identify candidate relationships, suggest classifications, compare standards, detect inconsistencies, and generate preliminary mappings between frameworks. Large collections of existing standards, controls, guidance documents, and implementation practices can therefore serve as input material for ontology development rather than requiring manual modelling from scratch.

Human experts remain essential because ontology engineering is ultimately a process of consensus creation and governance. Nevertheless, AI significantly reduces the effort required to transform document-based knowledge into machine-readable security models.

This observation naturally leads to the next topic. If the technical barriers to ontology creation are decreasing and if machine-readable standards can support AI, automation, and multi-compliance, the remaining challenge is no longer purely technological. It is organizational. Successful adoption depends upon collaboration, agreement on terminology, shared ownership, and a collective understanding of the domain being represented.

For this reason, the next chapter examines an often-overlooked aspect of ontology engineering: the role of collaboration and consensus in building a sustainable industry knowledge framework.



5. Collaboration and Consensus

The previous chapter examined how machine-readable knowledge, automation, reasoning engines, and artificial intelligence can increase the value of a security ontology. However, technology alone cannot create a shared understanding of cybersecurity. Before machines can reason over a domain, humans must first agree on how that domain should be described.

This observation is often overlooked. The most difficult challenge in ontology development is frequently not technical implementation, data modelling, or AI-assisted knowledge extraction. Rather, it is achieving agreement on concepts, definitions, responsibilities, and relationships. Ontologies are ultimately formal representations of human understanding. As a result, successful ontology engineering is as much a social activity as it is a technical one.

5.1

Shared conceptualization

A widely accepted definition describes an ontology as a formal specification of a shared conceptualization. The emphasis on *shared* is important.

Security professionals, architects, auditors, suppliers, regulators, and business leaders often use the same words while attaching different meanings to them. Terms such as *service*, *risk*, *control*, *supplier*, *asset*, *ownership*, or even *compliance* may be interpreted differently depending on professional background, company culture, or regulatory context.



This ambiguity becomes increasingly problematic as organizations attempt to automate governance processes, integrate multiple standards, and enable machine-assisted decision-making. Machines can process only what has been explicitly defined. Consequently, the value of an ontology depends directly on the degree to which stakeholders agree on the concepts and relationships it contains.

The development of a cybersecurity ontology therefore creates an opportunity to establish a common language across the industry. It transforms implicit assumptions into explicit definitions and provides a foundation upon which collaboration, automation, and interoperability can be built.

5.2

Industry collaboration

This challenge aligns closely with the mission of the Zero Outage Industry Standard (ZOIS).

Unlike an internal corporate framework, ZOIS is an industry association that brings together organizations with different perspectives, operational realities, and areas of expertise. Its purpose is not simply to publish requirements, but also to facilitate collaboration and promote a common understanding of secure and reliable service delivery.

An ontology could provide a focal point for this collaboration.



Rather than viewing standards as static documents that are periodically updated, participants could collaboratively contribute to a living knowledge model describing the domain itself. Concepts, controls, technologies, supplier relationships, architectural patterns, and compliance mappings could evolve within a shared framework while maintaining consistency and traceability.

Such an approach would provide more than a technical artifact. It would provide direction.

The ontology would express not only what the industry currently knows, but also how the community understands the relationships between concepts, where consensus exists, and where further discussion is required. In this sense, ontology development becomes a mechanism for coordinating expertise and aligning the collective intent of the association.

The resulting knowledge model could then support multiple outputs simultaneously: standards, guidance documents, compliance mappings, implementation tools, training materials, and AI-assisted services.

5.3

Why human consensus remains critical

The recent advances in artificial intelligence may create the impression that ontology development can eventually be delegated entirely to machines. In practice, the opposite is true.

AI can accelerate many activities associated with ontology engineering. It can extract concepts from documents, identify candidate relationships, compare standards, suggest mappings, and detect inconsistencies. As discussed in the previous chapter, these capabilities can significantly reduce the cost and effort of ontology development.


Yet AI does not carry responsibility.



The decisions encoded within an ontology ultimately influence compliance determinations, architectural choices, audits, governance processes, and business investments. These decisions have legal, financial, operational, and strategic consequences. Responsibility for such outcomes remains with organizations and the individuals who govern them.

Consequently, the essential function of AI should be viewed as assistance rather than authority.

Human experts remain responsible for:



An ontology therefore represents not machine-generated knowledge but human consensus captured in a machine-readable form.

This distinction is important because the long-term value of an industry ontology depends less on the sophistication of the technology and more on the credibility of the agreement behind it.



5.4

A strategic opportunity for ZOIS

The cybersecurity industry already contains several highly successful knowledge resources with ontology-like characteristics. MITRE D3FEND⁹, for example, explicitly publishes its knowledge model in OWL and RDF formats and describes itself as a cybersecurity knowledge graph¹⁰ supported by ontology resources.

Similarly, NIST research increasingly explores ontology-based approaches for integrating cybersecurity knowledge sources such as ATT&CK, D3FEND, CWE, and CVE into unified semantic models to improve interoperability and threat analysis¹¹.

However, these initiatives focus primarily on cyber threats, attack techniques, vulnerabilities, and defensive capabilities.

They are not industry-wide security governance standards comparable to ESARIS or the ZOIS Security Pillar. This distinction creates an opportunity.

To date, there is no widely adopted industry security standard whose primary representation is a comprehensive ontology supporting governance, architecture, supply-chain security, multi-compliance, and operational decision-making. Existing standards are generally delivered as structured documents, while ontology-based approaches remain complementary initiatives¹².

For ZOIS, this presents two possible strategic interpretations:



If ontology-driven standards become the future direction of cybersecurity governance, ZOIS should participate early to remain relevant and interoperable.



If such standards are not yet established, ZOIS has an opportunity to differentiate itself by becoming one of the first industry associations to provide a machine-readable semantic representation of its knowledge framework.

9 <https://d3fend.mitre.org/resources/ontology/>

10 <https://d3fend.mitre.org/>

11 <https://www.nist.gov/publications/design-and-application-unified-ontology-cyber-security>

12 <https://github.com/Cyber-Ontology-Foundry>



In either case, developing an ontology is not merely a technical exercise. It strengthens the ability of ZOIS to support multi-compliance, automation, AI integration, and supply-chain collaboration while creating a durable foundation for future growth.

The natural next question is therefore not whether such a vision is possible, but how it can be realized. The following chapter addresses this challenge by outlining a practical roadmap for transforming today's document-centric standards into an ontology-enabled security knowledge framework.




6. Roadmap

The previous chapter argued that the success of an ontology initiative depends as much on collaboration and consensus as it does on technology. If ZOIS decides to pursue such a path, the challenge is no longer conceptual but practical: how can the existing body of knowledge be transformed into an operational capability that supports multi-compliance, automation, and supply-chain security?

The answer is not a single project or technology. Rather, it is a sequence of incremental steps that build upon existing assets while progressively increasing their digital usability and business value. Fortunately, many of the underlying technologies and practices already exist, making the journey significantly more achievable than it would have been only a few years ago.

Figure 3 - Roadmap





6.1

Adopting machine-readable standards through OSCAL

The first step is to transform existing standard content into machine-readable formats.

A natural starting point is the Open Security Controls Assessment Language (OSCAL), an initiative led by NIST to represent security controls, assessments, implementation statements, and compliance artefacts in standardised machine-readable formats such as XML, JSON, and YAML. OSCAL was specifically designed to support interoperability between security and compliance tools while reducing the cost of exchanging and reusing control information.

For ZOIS, adopting OSCAL or a compatible representation offers several advantages:



Alignment with an internationally recognised approach



Easier integration with future tooling



Reuse of existing implementations, methodologies, and best practices



Compatibility with existing compliance and GRC ecosystems



Reduced effort when exchanging information with customers, suppliers, and regulators

Importantly, this is not an attempt to invent a new format. It is an opportunity to build upon an established foundation already supported by industry and government initiatives. By leveraging existing standards wherever possible, the transition toward machine-readable knowledge can begin with relatively limited disruption while creating long-term interoperability benefits.



6.2

Relationship modelling

Machine-readable controls alone do not constitute an ontology. The defining step of the roadmap is the modelling of relationships between concepts.

This activity represents the core deliverable of the proposed initiative. While OSCAL and similar technologies provide a structured representation of controls and requirements, relationship modelling introduces the semantic layer that transforms isolated elements into a connected knowledge system.

At this stage concepts such as:



The resulting model becomes the foundation for every capability discussed in the previous chapters. Multi-compliance mappings, AI-assisted analysis, reasoning engines, dependency analysis, change impact assessments, and future Digital Twin applications all depend on the existence of these relationships.



This is also the stage where important design decisions must be made. Choices regarding modelling approaches, engineering frameworks, governance structures, ontology languages, and quality assurance practices will significantly influence the long-term maintainability and usefulness of the ontology.

For this reason, the next chapter focuses specifically on the methodology and design principles that should guide ontology development.



6.3


User interfaces and accessibility

Even the most sophisticated ontology has limited value if it remains accessible only to ontology engineers.

Most users will never interact directly with OWL files, RDF triples, semantic databases, or machine-readable schemas. Instead, they require intuitive interfaces that allow them to benefit from the underlying knowledge without understanding its implementation details.

Consequently, the roadmap should include the development of presentation and interaction layers tailored to different audiences.

Examples include:



These interfaces serve an important purpose beyond simple usability. They create visibility and understanding. Stakeholders can see how concepts connect, explore the impact of decisions, and appreciate the value of ontology without needing expertise in semantic technologies.

In many respects, this layer plays a role similar to Harry Beck's Underground map discussed earlier in the paper. The complexity remains underneath, but the user experiences a navigable and approachable representation of the system.



6.4

Operational adoption

The final step of the roadmap is the most important. An ontology becomes valuable only when it is used. At this stage the ontology evolves from an industry knowledge model into an operational tool.

Organizations can adopt the shared ZOIS ontology and extend it to reflect their own environments, technologies, governance structures, and regulatory obligations.

This capability should not be viewed as fragmentation. Quite the opposite.

As a result:



organizations retain flexibility;



company-specific knowledge can be incorporated;



local requirements can be represented;



interoperability between organizations can be preserved.

This is especially important in supply-chain environments, where multiple organizations must collaborate while maintaining independent governance structures. A shared ontological foundation creates the possibility for meaningful communication between security models, compliance mappings, and governance processes across organizational boundaries.

From a ZOIS perspective, this objective aligns directly with the association's overarching mission. Improving security along the entire supply chain fundamentally requires a shared understanding of how security concepts relate to one another across organizations.



An ontology offers a mechanism for capturing and maintaining that shared understanding while still allowing individual participants to evolve according to their own business needs.

The roadmap outlined above is intentionally practical. It begins with machine-readable standards, progresses through relationship modelling and user accessibility, and ultimately leads to operational adoption. However, before such a roadmap can be executed, the methodological foundations must be established. The next chapter therefore examines the design principles and engineering decisions that should guide the development of a ZOIS ontology.



7. Methodology and Design Principles

The previous chapter outlined a practical roadmap for developing a ZOIS ontology. Before such an initiative can be executed, a set of design principles is required to guide development and ensure that the outcome remains useful, maintainable, and aligned with the objectives of the association.

The recommendations presented below build on established ontology engineering practices while remaining focused on the practical needs of ZOIS and its members. The intention is not to pursue academic perfection, but to develop an industry asset that can support multi-compliance, automation, and supply-chain security over the long term.

7.1

Hybrid approach

The proposed approach combines top-down design with bottom-up knowledge extraction.

The top-down perspective ensures that the ontology is built upon stable concepts and consistent definitions. The bottom-up perspective ensures that the model reflects the realities already captured within ESARIS, ZOIS standards, operational practices, and existing security documentation.

This combination allows the ontology to remain both structured and practical. Rather than developing a theoretical model disconnected from reality, or merely digitizing existing documents without improving them, the hybrid approach balances long-term consistency with operational relevance.



7.2

Enterprise-wide scope

The ontology should be developed with an enterprise-wide perspective.

While the initial implementation will focus on the Security Pillar and the ESARIS taxonomy, the model should accommodate future connections to business processes, IT architecture, operational activities, supplier relationships, and other ZOIS workstreams.

The objective is not to create an isolated compliance tool, but a knowledge structure capable of supporting broader organizational decision-making. This wider perspective also aligns with the supply-chain focus of ZOIS, where security cannot be considered independently from the systems and organizations it supports.

7.3

Collaborative development

The ontology should be developed collaboratively.

One of the underlying themes of this paper is that security knowledge is created through industry consensus. The same principle should apply to the ontology itself. ZOIS is not a single organization imposing a proprietary model, but an association bringing together expertise from multiple participants.

The development process should therefore encourage contributions from members, facilitate discussion around definitions and relationships, and maintain transparency regarding decisions and changes.

This approach not only improves the quality of the ontology but also increases acceptance and adoption among organizations that will ultimately use it.



7.4

Open and industry-supported technology foundations

The ontology should use an internationally recognized representation format capable of expressing detailed relationships and supporting future automation, reasoning, and integration requirements.

At the same time, technical choices should remain subordinate to business objectives. The purpose of the ontology is not to maximize technical sophistication, but to provide a reliable foundation for multi-compliance, decision support, interoperability, and future AI-assisted capabilities.

Choosing a mature and widely supported ontology technology ensures that the model remains portable, reusable, and compatible with future tools and platforms.

7.5

Structured development with agile delivery

The development process should combine a structured governance framework with agile working practices.

A structured framework provides clarity regarding responsibilities, quality expectations, documentation, and lifecycle management. Agile practices provide the flexibility to incorporate feedback, refine concepts, and respond to evolving requirements.

This balance allows the ontology to grow incrementally while maintaining overall coherence. Early deliverables can provide value to stakeholders, while the model continues to mature through successive iterations.



7.6

Continuous quality assurance

The ontology should be treated as a strategic asset and maintained accordingly.

Quality assurance should not be limited to technical validation. Equally important are expert reviews, consistency checks, verification of relationships, and periodic assessments of whether the model continues to reflect operational reality.

A combination of automated validation and regular stakeholder review provides an effective mechanism for ensuring accuracy and trustworthiness. As the ontology grows and becomes integrated into business processes, maintaining this quality discipline will become increasingly important.

Closing thoughts

The methodology proposed here intentionally favors practicality over complexity. A hybrid approach, enterprise-wide scope, collaborative governance, mature technology choices, structured development, and continuous quality assurance together provide a solid foundation for the initiative.

Most importantly, these principles reinforce a recurring message throughout this paper: the ontology is not the objective. The objective is to improve how organizations understand, manage, and collaborate on security across increasingly complex digital ecosystems. The ontology is simply the enabling mechanism that makes this possible.



8. Conclusion: Ontology as an Enabler of Resilient Multi-Compliance Capability

Cybersecurity is becoming increasingly complex. Organizations must navigate expanding regulatory obligations, growing supply-chain dependencies, evolving technologies, and rising expectations from customers, regulators, and business leaders. As a result, the challenge is no longer simply implementing security controls. The challenge is understanding how controls, risks, assets, services, suppliers, and compliance obligations relate to one another across an increasingly interconnected ecosystem.

This paper has argued that ontology technology offers a promising response to this challenge.

An ontology is not a replacement for standards, frameworks, governance processes, or human expertise. Rather, it provides a means of connecting these elements through a shared, machine-readable understanding of the domain.

ZOIS Opinion Paper

By making relationships explicit, ontologies enable organizations to navigate complexity more effectively, support multiple compliance frameworks simultaneously, improve decision-making, and create new opportunities for automation and AI-assisted analysis.

For ZOIS, the importance of this development extends beyond technology. The vision of strengthening security across supply chains depends fundamentally on collaboration between organizations. Supply-chain security cannot be achieved through isolated efforts. It requires a common understanding of security concepts, responsibilities, dependencies, and expectations that can be shared across organizational boundaries. An ontology provides a mechanism for building and maintaining that shared understanding while allowing individual organizations to adapt the model to their own environments and requirements.



The opportunity is therefore strategic. A machine-readable semantic representation of the ZOIS Security Pillar could become an important enabler for multi-compliance management, cross-organizational interoperability, AI-assisted governance, and ultimately more resilient digital ecosystems. Whether this vision is realized through incremental evolution or more ambitious transformation, the direction is consistent with the broader trends affecting cybersecurity, compliance, and digital governance.

Most importantly, this paper hopes to have conveyed that the true value of an ontology lies not in the technology itself, but in the collaboration it enables. The ontology is merely the tool; the real asset is the collective knowledge, experience, and consensus of the community that develops and maintains it.

If ZOIS can successfully bring together its members around a shared conceptualization of cybersecurity and supply-chain security, then an ontology may become far more than a knowledge model. It may become a foundation for stronger cooperation, more effective compliance management, and more resilient supply chains across the industries that ZOIS serves.



The author likes to thank his colleagues Erdélyi Judit, Thomas Hecht, and Márk Potkovác, for discussing the foundations of the ontology which considerably helped to elaborate it.