Finansuoja Europos Sąjunga
NextGenerationEU

Mykolo Romerio universitetas

NAUJOS KARTOS LIETUVA

Project „Misijomis grįstų mokslo ir inovacijų programų įgyvendinimas“ (Project Nr. 02-002-P-0001) report

Project name: ***CTI-BALANCED***

Responsible/Implementation partner (-s): ***MRU***

Action result: ***Integrated Privacy/GDPR compliance and cybersecurity compliance model***

Finansuoja Europos Sąjunga
NextGenerationEU

Mykolo Romerio universitetas

NAUJOS KARTOS LIETUVA

Editor
<<Main Editor>>

Contributors
Evaldas Bruze (MRU)
Giedrė Sabaliauskaitė (MRU)
R. Andrew Paskauskas (MRU)
Tomas Lavišius (MRU)

Reviewers
Evaldas Bružė (MRU)
Giedrė Sabaliauskaitė (MRU)
R. Andrew Paskauskas (MRU)
Tomas Lavišius (MRU)

Disclaimer

The information in this document is provided “as is”, and no guarantee or warranty is given that the information is fit for any particular purpose. The content of this document reflects only the author`s view. The users use the information at their sole risk and liability.

Executive Summary

This document presents the Privacy / GDPR Compliance and Cybersecurity Compliance Model developed by Mykolas Romeris University (MRU) as part of the CTI-BALANCED project. It directly fulfils the project deliverable requirement: "Sukurti privatumo / GDPR atitikties + kibernetinio saugumo atitikties modelį" (Create a privacy / GDPR compliance + cybersecurity compliance model). The model is grounded in a formally structured OWL/Turtle ontology, soc_cti_framework.ttl (namespace: https://purl.org/soc-cti-framework#) and its extended variant ircc-framework, implemented and validated in the Stardog Cloud knowledge graph environment.

The ontological framework formalises the EU regulatory structure relevant to Security Operations Centres (SOCs) and Cyber Threat Intelligence (CTI) sharing. It integrates 12 EU regulatory instruments (including NIS2, DORA, GDPR, AI Act, CRA, eIDAS), 20 requirements, 10 incident types, 8 reportable attributes, 3 communication standards (STIX, TAXII, MISP), 5 reporting timeframes, 10 penalties, and 3 accreditation requirements, all connected through a formal ontological structure with object properties that enable machine-readable compliance reasoning and SPARQL-based traceability chains.

The model comprises three layers. The core layer (Chapter 1) establishes the basic and granular ontological schemas. The GDPR and privacy compliance layer (Chapter 2) provides SPARQL-validated compliance chains across the full GDPR breach lifecycle, including breach classification, notification paths (Article 33 supervisory authority vs. Article 34 data subjects), lawful basis for CTI processing, and the Article 34 waiver logic for encryption and pseudonymisation. The cybersecurity compliance layer (Chapter 3) extends the framework to the MITRE ATT&CK for ICS threat matrix, producing 71 fully grounded regulatory traceability chains across 12 ICS attack tactics and 4 regulations (GDPR, NIS2, AI Act, CRA).

Table of Contents

List of Figures

No table of figures entries found.

List of Tables

No table of figures entries found.

Introduction

The CTI-BALANCED project addresses a fundamental challenge in collective cybersecurity: cyber threat intelligence sharing at the sectorial and national level takes place in a complex regulatory environment where privacy law, cybersecurity regulation, and data-sharing obligations intersect in ways that are rarely modelled formally. Compliance decisions, what data can be shared, with whom, on what legal basis, under what safeguards, are currently made through manual legal analysis, producing inconsistent and non-machine-readable outcomes.

This document presents MRU's response to this challenge: an ontological compliance model that formally represents the regulatory structure, enables SPARQL-based compliance querying, and produces end-to-end traceability chains from regulatory obligation through organisational requirement, security control, threat technique, and attack tactic. The model is designed to be reusable, extensible, and applicable to the CTI-BALANCED platform.

The model integrates four research artefacts developed sequentially during the project: (1) the initial soc_cti_framework.ttl ontology with basic and granular GDPR layers; (2) the quantified lightweight version with full SPARQL validation results; (3) the extended ircc-framework with MITRE ATT&CK for ICS integration; and (4) the scenario-based demonstration of CTI sharing traceability. Together they constitute a complete, formally validated privacy and cybersecurity compliance model.

1.1. DOA, SOTA Update [if needed]

1.2. Context Update [if needed]

2. Core Ontological Framework: Design and Architecture

The foundation of the compliance model is a standalone OWL/Turtle ontology, soc_cti_framework.ttl, with namespace https://purl.org/soc-cti-framework#. This ontology formalises the SOC regulatory structure from a data protection and cybersecurity compliance perspective, enabling machine-readable compliance analysis and SPARQL-based querying across multiple regulatory regimes.

2.1. Basic Schema

The basic schema establishes the core structural relationships of the SOC regulatory framework. Regulation sits at the apex and connects outward through eight primary object properties to the following class types: RequirementType, IncidentType, ReportableAttribute, CommunicationStandard, ReportingTimeframe, Penalty, AccreditationRequirement, and ResponsibleEntity. Each class type has both a class-level definition and concrete instances.

At the instance level, the basic schema is populated with 12 EU regulatory instruments: NIS2 (Directive (EU) 2022/2555), DORA (Regulation (EU) 2022/2554), CER Directive (Directive (EU) 2022/2557), eIDAS (Regulation (EU) No 910/2014), AI Act (Regulation (EU) 2024/1689), GDPR (Regulation (EU) 2016/679), CSA (Regulation (EU) 2019/881), NISD (original NIS Directive), DSA (Digital Services Act), ECRIS-TCN, ePrivacy Directive, and CRA (Cyber Resilience Act). The schema captures the core structural relationships: which regulations impose which requirements, which incident types trigger reporting obligations, which reportable attributes must be communicated in which formats, and which penalties apply to which requirement failures.

2.2. Granular GDPR Schema

The granular schema extends the basic structure with a detailed GDPR decomposition designed to represent the fine-grained logic of GDPR breach handling within a SOC/CTI context. This extension was developed to test whether the ontology could represent the kind of compliance logic actually required for operational decision-making in security incident scenarios involving personal data.

The granular layer adds seven new classes to the ontology:

- BreachType: classification of personal data breaches by their information security property: AvailabilityBreach, ConfidentialityBreach, IntegrityBreach.
- NotificationPath: the applicable notification obligation: Article33 (supervisory authority notification within 72 hours) vs. Article34 (direct notification to data subjects when high risk exists).
- PenaltyTier: the two-tier GDPR penalty structure: Tier1 (up to €10M or 2% global annual turnover) vs. Tier2 (up to €20M or 4% global annual turnover).
- LawfulBasis: the Article 6 GDPR processing grounds applicable to SOC/CTI operations: including legitimate interest (Article 6(1)(f)), legal obligation (Article 6(1)(c)), and public task (Article 6(1)(e)).
- DataSubjectCategory: categories of individuals whose data may be involved in CTI incidents: employees, customers, third parties, threat actors.
- DataCategory: specific personal data categories relevant to SOC operations: IP address, traffic data, hostname, email address, user credentials, authentication logs.
- Safeguard: technical measures that may affect notification obligations and compliance positioning: Encryption, Pseudonymisation, AccessControl, AuditLogging, IntegrityControl, BackupAndRecovery.

This extended architecture allows the ontology to represent the full granular logic of GDPR breach handling: a personal data breach is classified by its type, which determines the notification paths triggered, the data categories and data subjects involved, the safeguards required, and the penalty tier that applies. The instance layer beneath this granular schema includes 3 breach types, 6 data categories (IP address, traffic data, hostname, email address, user credentials, authentication logs), 4 data subject categories, 2 notification paths, 3 lawful bases, 6 safeguards, and 2 penalty tiers, all formally connected through object properties.

2.3. Object Property Structure

The ontology's expressive power derives from its object property structure. Key properties in the basic layer include: imposesRequirement (Regulation → RequirementType), triggersReporting (IncidentType → ReportableAttribute), requiresFormat (ReportableAttribute → CommunicationStandard), hasTimeframe (IncidentType → ReportingTimeframe), and appliesPenalty (Regulation → Penalty).

The granular GDPR layer adds: classifiedAs (PersonalDataBreach → BreachType), triggersNotification (BreachType → NotificationPath), involvesDataCategory (PersonalDataBreach → DataCategory), affectsDataSubject (PersonalDataBreach → DataSubjectCategory), requiresSafeguard (BreachType → Safeguard), carriesPenaltyTier (PersonalDataBreach → PenaltyTier), requiresLawfulBasis (RequirementType → LawfulBasis), and mayWaiveNotification (Safeguard → NotificationPath). This last property encodes the Article 34(3)(a) logic: if a SOC has implemented encryption or pseudonymisation on the affected data, it may be relieved of the obligation to notify data subjects individually.

2.4. Implementation and Validation Environment

The ontology is implemented in OWL/Turtle format and deployed in the Stardog Cloud knowledge graph environment. Stardog provides SPARQL 1.1 query support, OWL reasoning, and a visual graph schema explorer that was used to verify both the basic and granular schema layouts. All SPARQL queries described in this document were executed successfully in Stardog and validated against the instantiated ontology. The ontological artefacts and SPARQL query sets are available for transfer to project partners and for integration with the CTI-BALANCED platform.

3. GDPR and Privacy Compliance Model

The GDPR and privacy compliance layer operationalises the theoretical ontological structure into executable compliance queries. This chapter presents the quantified coverage of the lightweight ontology, the three foundational SPARQL queries that validate the basic schema, and the three advanced SPARQL queries that validate the granular GDPR layer, including the significant Article 34 waiver query.

3.1. Quantified Coverage of the Lightweight Ontology

The lightweight version of the ontology (soc_cti_framework.ttl, basic layer) covers the following instance populations:

- 12 regulations (NIS2, DORA, CER, eIDAS, AI Act, GDPR, CSA, NISD, DSA, ECRIS-TCN, ePrivacy, CRA)
- 20 requirements (RequirementType instances covering incident reporting, cybersecurity framework requirements, threat intelligence sharing, risk assessment, and further specialised obligations)
- 10 incident types
- 8 reportable attributes
- 3 communication standards: STIX, TAXII, MISP
- 5 reporting timeframes
- 10 penalties
- 3 accreditation requirements
- 2 responsible entity types

All 70+ class instances are interconnected through the object property structure described in Chapter 1, enabling SPARQL traversal across any combination of regulatory dimension.

3.2. Basic Schema Queries: Cross-Regulatory Analysis

Query 1: Cross-Regulatory Requirement Overlap

This query identifies which requirements are imposed by the greatest number of regulations, confirming that incident reporting and cybersecurity framework requirements are the most pervasive obligations across the EU regulatory structure. The full result set (20 requirements) is shown in Table 2.1.

Requirement	Regulations Imposing It (count)
IncidentReporting	7
CybersecurityFrameworkReq	7
ThreatIntelligenceSharing	5
RiskAssessment	4
IncidentResponse	4
CrossBorderCollaboration	4
CertificationCompliance	2
IncidentAnalysis	2
TransparencyReq	1
LawEnforcementCollaboration	1
StrongAuthentication	1
SecurityOfProcessing	1
ConfidentialityOfCommunications	1
DPOAppointment	1
DPIA	1
DataBreachNotification	1
DataProtectionByDesign	1
ThreatManagement	1
CriticalInfraProtection	1
OperationalResilience	1

Query 2: Full Reporting Chain

This query traces the complete path from regulation → incident type → reportable attributes → communication format, producing structured end-to-end reporting chains for each regulation-incident combination. The query confirms that CTI artefacts (indicators of compromise, threat actor profiles, malware signatures) must be communicated in STIX or TAXII format under NIS2 and DORA, while MISP is the preferred platform for ISAC-level sharing. The full result set is available from the Stardog SPARQL endpoint.

Query 3: Penalty Exposure per Requirement

This query links regulations to requirements to specific penalty amounts, surfacing the financial liability map for each compliance obligation. Results confirm that GDPR imposes the highest penalty ceiling (€20M / 4% global turnover for Tier 2 violations) and that NIS2 and DORA each provide substantial penalty exposure for failure to report incidents and maintain cybersecurity risk management frameworks. The penalty exposure map provides the compliance prioritisation logic for the CTI-BALANCED platform's governance model.

3.3. Granular GDPR Layer: Extended Compliance Queries

The granular GDPR extension adds seven new classes and a substantially extended instance set to the ontology, enabling three advanced compliance queries relevant to SOC/CTI operational decision-making.

Query 4 (d): Breach-to-Penalty Chain

The full breach-to-penalty chain traverses: PersonalDataBreach → BreachType → DataCategory → DataSubject → NotificationPath → Timeframe → Safeguard → PenaltyTier → MaxAmount. Execution of this query against the granular ontology layer produced 120 fully grounded results spanning AvailabilityBreach, ConfidentialityBreach, and IntegrityBreach. Table 2.2 shows the schema of this query chain (column headers); the full 120-row result set is available from the Stardog endpoint.

breachType	dataCategory	dataSubject	notification	timeframe	safeguard	penaltyTier	maxAmount
(120-row full result set available from Stardog SPARQL endpoint)							

Query 5 (e): Lawful Basis Chain

This query surfaces which SOC requirements need which GDPR Article 6 legal grounds, and how those processing grounds connect upstream to the regulations that impose the relevant obligations. Key findings: threat intelligence sharing under NIS2 and DORA can rely on legitimate interest (Article 6(1)(f)) or legal obligation (Article 6(1)(c)) depending on whether the sharing organisation is subject to a binding sectoral sharing obligation. Purely voluntary CTI sharing between private entities requires a well-documented legitimate interest assessment. SOC monitoring of employee communications for security purposes may engage public task (Article 6(1)(e)) for public-sector entities or legitimate interest for private-sector entities.

Query 6 (f): Article 34 Waiver Query

This is analytically the most significant query because it surfaces the practical compliance decision that a SOC faces after a personal data breach involving CTI data: if encryption and pseudonymisation were applied to the affected data, Article 34(3)(a) GDPR may relieve the SOC of the obligation to notify data subjects individually, saving significant operational effort and reputational exposure, while still requiring supervisory authority notification under Article 33.

The query returned two rows, both showing that only the Confidentiality Breach triggers the Article 34 data subject notification path, and that the two safeguards capable of activating the waiver are Encryption and Pseudonymisation. Integrity and Availability breaches do not appear because they are not linked to the DataSubjectNotification path in the ontology, reflecting the legal logic that confidentiality breaches are the primary scenario where data subjects face high risk from unauthorised disclosure. Table 2.3 shows the Article 34 waiver result.

breachType	dataCategory	safeguard	notificationWaived
ConfidentialityBreach	EmailAddress / UserCredentials	Encryption	DataSubjectNotification (Art.34)
ConfidentialityBreach	EmailAddress / UserCredentials	Pseudonymisation	DataSubjectNotification (Art.34)

This finding is applicable to the CTI-BALANCED platform design: any CTI sharing architecture that applies encryption and pseudonymisation to all shared network identifiers and personal data will benefit from the Article 34(3)(a) waiver in the event of a platform-level breach, materially reducing the notification burden on participating organisations. This makes encryption and pseudonymisation not merely good practice but compliance-optimal design decisions with quantifiable regulatory benefit.

3.4. Privacy Compliance Model: Summary

The GDPR and privacy compliance layer of the ontological model establishes three practical compliance outputs for the CTI-BALANCED platform:

- Regulatory coverage map: which of the 12 regulatory instruments impose which of the 20 requirement types, enabling compliance prioritisation by regulatory density.
- Breach handling decision support: the full breach-to-penalty chain enables automated compliance triage for any personal data breach scenario encountered in CTI sharing operations.
- Notification minimisation strategy: the Article 34 waiver query formalises the compliance benefit of encryption and pseudonymisation as platform-level privacy-by-design controls.

4. Cybersecurity Compliance Model: ICS Extension and Regulatory Traceability

The cybersecurity compliance layer extends the ontological framework from the GDPR and data protection domain into the cybersecurity engineering domain, enabling formal traceability from regulatory obligations through organisational security requirements, technical controls, specific threat techniques, and attack tactics. This extension directly bridges the gap between legal compliance analysis and operational cybersecurity defence, the core challenge in CTI sharing for critical infrastructure.

4.1. Ontological Extension: The Tactic Class

To enable systematic integration of the MITRE ATT&CK for ICS matrix, a new class (Tactic) was introduced into the ontology, together with a corresponding object property: belongsToTactic (ThreatTechnique → Tactic). This addition allows the existing SecurityControl-to-ThreatTechnique mitigation structure to be extended upward to the tactic level.

Critically, each tactic was mapped to one or more NIST CSF 2.0 functions via the mapsToFunction property. This tactic-to-function mapping establishes the semantic bridge between the threat model (what adversaries do) and the regulatory model (what organisations are required to achieve). The result is a connected path across the full ontological graph:

Regulation → RegulatoryRequirement → SecurityControl → ThreatTechnique → Tactic → NIST CSF Function

This path makes it possible to query: "which regulations impose requirements that, when implemented as security controls, mitigate techniques belonging to a specific attack tactic?", producing the regulatory traceability chains that are the primary deliverable of this layer.

The ontology was further extended with six additional regulatory requirements: risk management, vulnerability handling, supply chain security, business continuity, access control, and security by design. These additions bring the cybersecurity compliance framework into direct coverage of NIS2 Article 21, DORA Article 16, CRA Article 13, and AI Act Article 15.

4.2. Table X1: Tactic-to-Function Mapping

Table 3.1 presents the mapping between the 12 MITRE ATT&CK for ICS tactics and NIST CSF 2.0 functions. This mapping constitutes the primary integration layer between the threat model and the regulatory model. Each row shows a tactic, the CSF function(s) it maps to, and the regulatory instruments that impose requirements directly relevant to that tactic via the mapped function.

Tactic (MITRE ATT&CK ICS)	**NIST CSF 2.0 Function(s)**	**Primary Regulatory Triggers**
Initial Access (TA0108)	Protect	NIS2 Access Control; CRA Security by Design
Execution (TA0104)	Protect, Detect	AI Act Security by Design; NIS2 Access Control; AI Act/GDPR Human Oversight
Persistence (TA0110)	Protect, Detect	CRA Vulnerability Handling; AI Act Security by Design; NIS2 Access Control
Privilege Escalation (TA0111)	Protect	NIS2 Access Control
Evasion (TA0103)	Detect	AI Act/GDPR Human Oversight
Discovery (TA0102)	Detect, Identify	NIS2/AI Act Risk Management
Lateral Movement (TA0109)	Protect, Detect	NIS2 Access Control
Collection (TA0100)	Detect, Protect	GDPR Data Minimization; AI Act/GDPR Human Oversight
Command and Control (TA0101)	Detect, Protect	GDPR Data Minimization

Tactic (MITRE ATT&CK ICS)	NIST CSF 2.0 Function(s)	Primary Regulatory Triggers
Inhibit Response Function (TA0107)	Respond, Recover, Detect	AI Act/GDPR Human Oversight; NIS2 Business Continuity
Impair Process Control (TA0106)	Detect, Respond	AI Act/GDPR Human Oversight
Impact (TA0105)	Respond, Recover	NIS2 Business Continuity; NIS2 Incident Reporting

4.3. Extended SPARQL Query and Traceability Chains

The SPARQL query that produces the full regulatory-to-threat traceability chains uses the following structure (ircc-framework namespace):

```
PREFIX irccf: <https://purl.org/ircc-framework#>
SELECT ?regulation ?requirement ?function ?tactic ?technique ?control
WHERE {
 ?regulation irccf:imposesRequirement ?requirement .
 ?requirement irccf:mapsToFunction ?function .
 ?requirement irccf:implementedBy ?control .
 ?control irccf:mitigates ?technique .
 ?technique irccf:belongsToTactic ?tactic .
 ?tactic irccf:mapsToFunction ?function .
}
ORDER BY ?tactic ?regulation
```

This query exploits the tactic layer to join the regulatory requirement graph to the threat technique graph through a common NIST CSF function variable. Unlike an asset-centred query, this approach ensures that only semantically coherent chains are returned, those where the regulatory function maps to the tactic function. Execution in Stardog produced 71 fully grounded traceability chains spanning all 12 ICS tactics and all four regulatory sources (GDPR, NIS2, AI Act, CRA).

Table 3.2 presents 17 representative fully grounded traceability chains selected to illustrate coverage across all 12 tactics.

Regulation	Requirement	Function	Tactic	Technique	Control
GDPR	Data Minimization	Protect	Collection	Screen Capture (T0852)	Data Confidentiality
GDPR	Data Minimization	Protect	Collection	Adversary-in-the-Middle (T0830)	Network Monitoring
GDPR	Data Minimization	Protect	Command and Control	Commonly Used Port (T0885)	Network Monitoring
NIS2	Risk Management	Identify	Discovery	Network Sniffing (T0842)	Network Monitoring
AI Act	Human Oversight	Detect	Evasion	Spoof Reporting Message (T0856)	Anomaly Detection
GDPR	Human Oversight	Detect	Evasion	Indicator Removal on Host (T0872)	Audit Logging

Regulation	Requirement	Function	Tactic	Technique	Control
NIS2	Access Control	Protect	Execution	Command-Line Interface (T0807)	Use Control Authorization
CRA	Vulnerability Handling	Protect	Execution	Scripting (T0853)	Input Validation
AI Act	Security by Design	Protect	Execution	Modify Controller Tasking (T0821)	Anomaly Detection
NIS2	Access Control	Protect	Initial Access	External Remote Services (T0822)	Access Control Enforcement
AI Act	Human Oversight	Detect	Impair Process Control	Unauthorized Command Message (T0855)	Anomaly Detection
NIS2	Business Continuity	Recover	Inhibit Response Function	Denial of Service (T0814)	Resource Availability
GDPR	Human Oversight	Detect	Inhibit Response Function	Alarm Suppression (T0878)	Anomaly Detection
NIS2	Access Control	Protect	Lateral Movement	Default Credentials (T0812)	Access Control Enforcement
NIS2	Access Control	Protect	Persistence	Valid Accounts (T0859)	Access Control Enforcement
CRA	Security by Design	Protect	Persistence	Module Firmware (T0839)	System Integrity
NIS2	Business Continuity	Recover	Impact	Loss of Availability (T0826)	Backup and Recovery

4.4. Table X3: Regulatory Density by Tactic

Table 3.3 summarises the distribution of all 71 fully grounded traceability chains across the 12 MITRE ATT&CK for ICS tactics and the 4 regulatory instruments, revealing differential regulatory density across phases of the ICS attack chain.

Tactic	Total Chains	Regulations Involved
Execution	17	AI Act, CRA, GDPR, NIS2
Persistence	14	AI Act, CRA, NIS2
Evasion	8	AI Act, GDPR
Collection	7	GDPR
Inhibit Response Function	6	AI Act, GDPR, NIS2
Impair Process Control	6	AI Act, GDPR
Discovery	4	AI Act, NIS2
Initial Access	3	AI Act, CRA, NIS2
Command and Control	2	GDPR
Lateral Movement	1	NIS2

Tactic	Total Chains	Regulations Involved
Privilege Escalation	1	NIS2
Impact	1	NIS2
Total	71	

4.5. CTI Sharing Scenario: Regulatory Traceability in Practice

The ontological framework was applied to a CTI sharing scenario complementing the ICS attack chain analysis. This scenario demonstrated that the traceability structure extends beyond compliance-to-threat mapping into the operational domain of incident reporting and intelligence sharing.

The instantiated traceability chains showed how NIS2 incident reporting requirements propagate through the Respond function to operational sharing protocols (STIX/TAXII), while GDPR Data Minimisation requirements constrain what personal data can be included in shared CTI artefacts. This creates a formally navigable tension: NIS2 incentivises comprehensive incident data sharing for collective cybersecurity, while GDPR requires minimisation of personal data in shared records.

The ontological representation makes this tension formally navigable. By linking the CTI platform to both the incident reporting requirement (via appliesTo: CTI-BALANCED Platform) and the data minimisation requirement (via appliesTo), the framework surfaces the specific design decision: pseudonymise IP addresses and user identifiers before sharing, retain original data for the supervisory authority notification, and document the pseudonymisation as a safeguard in the DPIA. This is the compliance-by-design output that the CTI-BALANCED platform requires.

The ontological structure also supports the integration of CTI sharing with the ATT&CK-based threat classification from Scenario 1. An ICS alert categorised as belonging to the Evasion tactic (e.g., Alarm Suppression T0878) can be traced to the GDPR/AI Act Human Oversight requirement via the Detect function, which then determines the applicable reporting obligation and the CTI sharing format. This cross-scenario integration is the most practically useful capability of the combined ontological framework.

4.6. Discussion: Contributions and Analytical Findings

The central contribution of this work is methodological: the demonstration that regulatory requirements, organisational cybersecurity frameworks, technical standards, threat classification models, and operational CTI sharing protocols can be integrated within a single formal ontological structure. This integration produces machine-readable, SPARQL-queryable compliance intelligence that is not available from any existing regulatory analysis tool or compliance framework.

The introduction of the Tactic class as an intermediary between individual threat techniques and NIST CSF 2.0 functions proved to be a structurally productive design choice. It enables the regulatory model to operate at the tactic level, which is the level at which CTI sharing communities typically reason about threats, rather than at the individual technique level, which is too granular for regulatory analysis.

4.7. Regulatory Asymmetry and Convergence Findings

Four significant findings emerge from the 71 traceability chains:

1) regulatory asymmetry across the ICS attack chain. Execution and Persistence account for 31 of the 71 chains (44%), indicating substantially denser regulatory coverage for early-phase attack prevention. The Impact tactic, where critical infrastructure damage occurs, has only 1 chain (NIS2 Business Continuity). This asymmetry indicates that the current EU regulatory structure provides substantially denser prescriptive coverage for preventive and protective measures than for response and recovery. This is a finding with direct policy implications for NIS2 implementation guidance.

2) regulatory convergence at the technique level. Techniques such as Modify Controller Tasking (T0821) and Change Operating Mode (T0858) are reached by multiple independent regulatory pathways:

- AI Act Security by Design → Protect
- GDPR Human Oversight → Detect

This convergence, in which distinct legislative instruments independently mandate controls for the same adversarial behaviour, is not codified within any existing cross-regulatory compliance tool. The ontological framework makes it formally explicit and queryable.

3) the Evasion tension. The Evasion tactic presents a structural tension within the regulatory compliance framework. Techniques such as Alarm Suppression (T0878), Spoof Reporting Message (T0856), and Indicator Removal on Host (T0872) map exclusively to the Detect function, and the regulatory coverage comes primarily from GDPR and AI Act human oversight requirements rather than from cybersecurity-specific legislation. This means that the primary regulatory driver for detecting evasion techniques in OT/ICS environments is data protection and AI governance law, a unexpected finding with significant implications for how compliance programmes prioritise human oversight controls.

4) single-control dependency risk. The Anomaly Detection Control carries the majority of Human Oversight traceability chains across four tactics (Evasion, Collection, Inhibit Response Function, Impair Process Control). This structural dependency means that a single control failure, or a deliberate attack that defeats anomaly detection, eliminates regulatory compliance coverage across 27 of the 71 chains. This finding supports the case for requiring multiple independent detection controls in NIS2 implementation for critical infrastructure.

4.8. Limitations

The following limitations of the current framework are acknowledged:

- The ontological framework captures a static representation of the regulatory structure. EU regulations and implementing acts change, and the ontology requires periodic maintenance to remain current.
- The scope of threat model integration is currently limited to MITRE ATT&CK for ICS. The Enterprise and Mobile matrices, relevant to many CTI sharing scenarios, are not yet integrated.
- The framework does not currently employ SHACL constraint validation to enforce structural or completeness constraints. SHACL layers are a priority for the next development phase.
- The tactic-to-function mapping involves interpretive judgments. Where tactics span multiple CSF functions, the assignment reflects the primary regulatory interface and may not capture all relevant secondary functions.
- The illustrative scenarios have not been validated in operational settings with critical infrastructure operators. Field validation is required before the framework is used for binding compliance decisions.

4.9. Future Work

The following development directions are identified:

- SHACL constraint layers: enforcement of ontological integrity and compliance completeness through shape-based constraint validation.
- Automated regulatory update ingestion: semi-automated methods for incorporating regulatory updates, including new legislation, implementing acts, and supervisory authority guidance.
- Cross-domain threat model integration: extension to MITRE ATT&CK Enterprise and Mobile matrices for non-ICS CTI sharing scenarios.
- Cross-scenario integration: formal linking of the ICS threat traceability (Scenario 1) with the CTI sharing compliance chains (Scenario 2), enabling ATT&CK-classified threat data to automatically trigger the applicable reporting chain and data minimisation logic.
- Sector-specific instantiation: application of the framework to all NIS2-defined critical sectors (energy, transport, health, digital infrastructure, public administration, banking, financial market infrastructure) to produce sector-tailored compliance intelligence.
- Operational validation: collaborative engagement with national cybersecurity authorities and critical infrastructure operators for field validation and practical deployment.

5. Summary and Conclusion

This document presents a complete, formally validated Privacy / GDPR Compliance and Cybersecurity Compliance Model, fulfilling the CTI-BALANCED project requirement "Sukurti privatumo / GDPR atitikties + kibernetinio saugumo atitikties modelį". The model is grounded in an OWL/Turtle ontology (soc_cti_framework.ttl / ircc-framework), deployed and validated in Stardog Cloud, covering 12 EU regulatory instruments, 20 requirements, the full GDPR breach lifecycle, and 71 fully grounded regulatory traceability chains across the MITRE ATT&CK for ICS matrix.

The model's three-layer architecture (core ontological framework, GDPR/privacy compliance layer, and cybersecurity compliance layer) provides a coherent, machine-readable compliance knowledge base for the CTI-BALANCED platform. The key compliance outputs are: the regulatory coverage map (which regulations impose which requirements with what penalty exposure), the breach handling decision support chain (full GDPR breach-to-penalty traceability), the Article 34 waiver logic (encryption and pseudonymisation as compliance-optimal design decisions), and the regulatory traceability from legal obligation to ICS threat technique (71 chains).

Four significant findings have policy implications beyond the CTI-BALANCED project: (1) regulatory asymmetry between preventive and recovery phases of ICS attacks indicates a gap in current EU regulatory coverage for response to critical infrastructure impact; (2) regulatory convergence at the technique level demonstrates that cross-regulation compliance analysis is essential for avoiding fragmented control implementation; (3) the Evasion-detection tension reveals that data protection law (GDPR/AI Act) is the primary regulatory driver for OT/ICS evasion detection, not cybersecurity-specific legislation; (4) single-control dependency risk across 27 of 71 chains supports mandatory multi-layer detection requirements in NIS2 implementation guidance. The ontological framework and all SPARQL query sets are available for partner transfer and integration into the CTI-BALANCED platform architecture.

List of Abbreviations

Abbreviation	Translation
CTI	Cyber Threat Intelligence
MRU	Mykolas Romeris University / Mykolo Romerio universitetas
SOC	Security Operations Centre
GDPR	General Data Protection Regulation (Regulation (EU) 2016/679)
NIS2	Network and Information Systems Security Directive 2 (Directive (EU) 2022/2555)
DORA	Digital Operational Resilience Act (Regulation (EU) 2022/2554)
CRA	Cyber Resilience Act
AI Act	Regulation (EU) 2024/1689
CER	Critical Entities Resilience Directive (Directive (EU) 2022/2557)
CSA	EU Cybersecurity Act (Regulation (EU) 2019/881)
eIDAS	Electronic Identification and Trust Services Regulation (Regulation (EU) 910/2014)
OWL	Web Ontology Language
TTL / Turtle	Terse RDF Triple Language — OWL serialisation format
SPARQL	SPARQL Protocol and RDF Query Language
NIST CSF	NIST Cybersecurity Framework 2.0
ICS	Industrial Control Systems
MITRE ATT&CK	MITRE Adversarial Tactics, Techniques and Common Knowledge framework
STIX	Structured Threat Information Expression
TAXII	Trusted Automated Exchange of Intelligence Information
MISP	Malware Information Sharing Platform
ISAC	Information Sharing and Analysis Centre
SHACL	Shapes Constraint Language
IEC 62443	Industrial Automation and Control Systems security standard series
DPIA	Data Protection Impact Assessment
PDAV	Poveikio duomenų apsaugai vertinimas (Lithuanian: DPIA)
TRL	Technology Readiness Level

Bibliography

soc_cti_framework.ttl ontology, namespace: https://purl.org/soc-cti-framework#. Implemented in Stardog Cloud. MRU CTI-BALANCED project artefact.

ircc-framework ontology, namespace: https://purl.org/ircc-framework#. Extended ICS version. MRU CTI-BALANCED project artefact.

Regulation (EU) 2016/679 of the European Parliament and of the Council (GDPR). Available at: https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng

Directive (EU) 2022/2555 (NIS2). Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32022L2555

Regulation (EU) 2022/2554 (DORA). Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32022R2554

Regulation (EU) 2024/1689 (EU AI Act). Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L_202401689

Directive (EU) 2022/2557 (CER Directive). Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32022L2557

NIST Cybersecurity Framework 2.0. NIST, 2024. Available at: https://www.nist.gov/cyberframework

MITRE ATT&CK for ICS. MITRE Corporation. Available at: https://attack.mitre.org/matrices/ics/

IEC 62443-3-3: Industrial Automation and Control Systems: System security requirements and security levels. IEC, 2013.

ENISA. (2025). NIS2 Technical Implementation Guidance. Available at: https://www.enisa.europa.eu/publications/nis2-technical-implementation-guidance

Bružė, E.; Paskauskas, R.A.; Matulytė, R. et al. Integration of Hybrid Threat Intelligence: The HiPSTer Ontological Methodology. Open Research Europe 2025. DOI: 10.12688/openreseurope.18929.1

W3C OWL 2 Web Ontology Language. W3C Recommendation, 2012. Available at: https://www.w3.org/TR/owl2-overview/

Stardog Knowledge Graph Platform. Available at: https://www.stardog.com/

FIRST CSIRT Services Framework v2.1. FIRST, 2023. Available at: https://www.first.org/education/csirt_service_framework