Ontology-driven Compliance Reasoning for the EU Cyber Resilience Act: A Proof-of-Concept

Jasmin Cosic
Information and Cyber Security Department
DEKRA SE
Stuttgart, Germany
jasmin.cosic@dekra.com

Miroslav, Baca
CEO and Head of R&D
Cyber-security, LTD
Zagreb, Croatia
ceo@cyber-security.hr

Abstract— **The Cyber Resilience Act (CRA) introduces new regulatory obligations for products with digital elements, creating significant challenges for scalable and explainable conformity assessment. This paper presents a proof-of-concept system for ontology-based pre-market conformity triage under the CRA, operationalizing the previously proposed OntoCRA-NS (Neuro-Symbolic Framework) reference architecture. The system combines structured evidence extraction, ontology population, and symbolic reasoning to infer regulatory admissibility outcomes based on a representative subset of CRA requirements. Ontological modelling and rule-based reasoning are used to evaluate the presence and adequacy of key conformity artefacts, such as technical documentation, software bill of materials, vulnerability status, and security update support. In the present prototype, evidence extraction is implemented using deterministic heuristics to validate the end-to-end reasoning pipeline. The system produces explainable triage outcomes, including compliant, non-compliant, and cases requiring human assessment. Evaluation is performed using curated product dossiers from the CURIUM project and an expert-defined reference labelling protocol to verify reasoning correctness and consistency. The results demonstrate the feasibility of ontology-driven regulatory triage and highlight the role of symbolic AI (Artificial Intelligence) in supporting transparent and human-centric regulatory decision-making..**

Keywords— Cyber Resilience Act, regulatory compliance, certification, neuro-symbolic AI, ontology reasoning, SWRL, evidence extraction, SBOM

I. INTRODUCTION

The EU Cyber Resilience Act (CRA) [1] introduces horizontal cybersecurity and vulnerability management requirements for products with digital elements across their lifecycle. Compliance must be demonstrated prior to market access through technical documentation and an EU Declaration of Conformity (DoC), while manufacturers must also manage vulnerabilities and maintain security support over time. In practice, conformity assessment remains largely manual, expert-driven, and difficult to scale across complex software supply chains.[1] Automated approaches based on machine learning and natural language processing can interpret unstructured documentation, yet they typically lack formal reasoning and transparent justification—capabilities required for legally defensible regulatory decision-making. Conversely, ontology-driven compliance reasoning provides deterministic and auditable inference but assumes structured input facts, which rarely exist in industrial dossier collections.

To bridge this gap, the OntoCRA-NS (Ontology based Neuro-Symbolic CRA assessment) [2] framework proposes a five-layer neuro-symbolic architecture combining neural evidence extraction with an OWL (Ontology Web Language) and SWRL (Semantic Web Rule Language) reasoning core executed by a description logic reasoner. This companion paper focuses on a practical validation: we operationalize OntoCRA-NS for pre-market conformity dossier triage and evaluate evidence extraction quality, end-to-end compliance decision correctness, and explanation traceability on six dossiers. From a Regulatory Technology (RegTech) perspective, CRA conformity assessment requires the transformation of legal obligations into machine-interpretable knowledge and the systematic evaluation of technical evidence against these obligations.

While the OntoCRA-NS architecture supports LLM-assisted (Large Language Model) evidence extraction as part of a neuro-symbolic design, the present proof-of-concept implementation employs a deterministic, rule-based extractor to ensure reproducibility and traceability. The integration and quantitative evaluation of LLM-based extraction is left for future work. A rule-based extractor was intentionally selected to align with regulatory expectations of determinism and auditability. Anyway, this paper makes the following contributions:

1. An end-to-end PoC (Proof of Concept) pipeline that transforms heterogeneous CRA technical documentation into structured Evidence JSON (JavaScript Object Notation), populates OntoCRA v3, and produces deterministic compliance outcomes with explanations.
2. An operationalization of conservative reasoning for CRA triage via an explicit Assessment Pending outcome when decisive evidence is missing.
3. A scenario-driven dossier dataset and evaluation protocol focusing on evidence correctness, decision agreement, and a safety metric (false compliant cases).
4. A traceability-oriented reporting workflow that links inferred decisions to rule premises and source evidence excerpts to support auditability.

Scope and intended use: OntoCRA-NS is designed as a decision-support tool for pre-market conformity dossier triage. It helps structure evidence, identify gaps, and generate explainable draft decisions under the modelled subset of CRA

requirements. Final conformity assessment remains a human and/or notified-body responsibility.

MATCH & CONTRIBUTION

This paper aligns with the thematic focus of ICE2026 on digital transformation, artificial intelligence, and governance by addressing the challenge of regulatory compliance in the context of the European Union Cyber Resilience Act (CRA). It contributes to the domains of Regulatory Technology (RegTech), cybersecurity assurance, and explainable AI by proposing a structured and transparent approach to compliance reasoning. The work fits within the conference scope by combining knowledge representation, semantic technologies, and neuro-symbolic AI to support decision-making in complex regulatory environments. The proposed approach enables traceable and explainable compliance assessments, which are critical for both industrial adoption and regulatory acceptance.

The main contributions of this paper are as follows:

• An ontology-driven compliance framework that formalizes CRA requirements and supports structured reasoning over regulatory artefacts;

• A neuro-symbolic pipeline integrating evidence extraction from technical documentation with symbolic reasoning based on Semantic Web Rule Language (SWRL);

• A triage mechanism that distinguishes between compliant, non-compliant, and assessment-pending cases to handle incomplete or uncertain evidence;

• A proof-of-concept implementation demonstrating the feasibility and practical applicability of the approach in realistic compliance scenarios.

By bridging unstructured documentation and formal compliance models, this work contributes to advancing automated and explainable regulatory assessment in cybersecurity.

II. ONTOCRA-NS FRAMEWORK RECAP

OntoCRA-NS is a modular neuro-symbolic compliance framework designed as a five-layer pipeline:

- Input Layer - gathers heterogeneous evidence sources,
- Neural Extraction Layer - interprets unstructured data,
- Ontology Adapter Layer - maps extracted facts into formal ontology instances,
- Symbolic Reasoning Layer - executes CRA-specific SWRL rules, and
- Output & Feedback Layer - generates explainable compliance reports and supports human validation. [2]

This separation of concerns provides flexibility: the neural components can evolve independently (e.g., LLM → fine-tuned model), while the ontology and rule base remain stable and transparent, ensuring consistent assessment procedures. It is important to mention that in the present proof-of-concept, the neural extraction layer is simulated using deterministic pattern-based extractors in order to isolate and validate the ontology population and reasoning layers. A key architectural element is the Evidence JSON contract between neural and symbolic components. Extracted compliance attributes (e.g., *SBOM – Software Bill of Materials presence, vulnerability handling process, known exploited vulnerability status, support duration*) are structured into JSON and mapped via Owlready2 to individuals and properties of the class *ProductWithDigitalElements*. OntoCRA v3 formalizes CRA Essential and Vulnerability Requirements using OWL classes and SWRL rules, executed by the Pellet reasoner, to infer deterministic outcomes such as *DoC_Granted, BannedFromMarket*, and optionally "*Assessment_Pending*" when evidence is insufficient.

III. RELATED WORK AND STATE OF THE ART

Regulatory compliance automation has been widely studied under the umbrella of RegTech [3], where the goal is to reduce manual effort and improve auditability through machine-interpretable representations of regulatory obligations. Butler characterizes RegTech as an enabling layer for digital compliance, emphasizing automation and traceability as key drivers for scalable governance. [4] Building on this direction, multiple works have operationalized compliance checking by encoding regulatory requirements into ontologies and executing inference with rule engines, demonstrating that formal knowledge models can support systematic, repeatable compliance analysis rather than ad-hoc checklist assessment.[5] Within ontology-driven compliance reasoning, SWRL [6] and SQWRL[7] based approaches are frequently used to convert normative requirements into executable constraints and assessment queries. For example, an SQWRL-based compliance prototype for regulated environments illustrates how a domain ontology combined with rule/query layers can produce deterministic compliance outputs and explanations, which is directly aligned with the explainability goal of CRA-oriented reasoning systems.[5] In parallel, the RegTech literature also includes OWL [9] based ontology approaches designed to support granular semantic querying and inference as part of compliance automation pipelines, reinforcing the feasibility of ontology-cantered architectures for real-world regulatory use.[4]

In cybersecurity, ontologies have increasingly been proposed to structure certification and assurance knowledge, enabling interoperable representations of security claims, controls, and assessment evidence. Recent work from the CUSTODES project context proposes an ontology tailored for cybersecurity certification of composite systems, explicitly motivating automation and efficiency gains in certification workflows—an adjacent problem to CRA conformity assessment and one that supports the broader relevance of ontology-driven certification reasoning.[11][12] Complementary surveys highlight a growing convergence between cybersecurity ontologies, knowledge graphs, and language models for structuring unstructured artifacts and supporting security analysis at scale, motivating hybrid neuro-symbolic pipelines where extraction feeds formal reasoning.[12]

A key evidence artifact emphasized in modern software security governance is the Software Bill of Materials (SBOM). The NTIA's "minimum elements" report established baseline expectations for SBOM content and has become a foundational reference for SBOM operationalization across ecosystems.[13] More recently, ENISA has published SBOM implementation guidance aimed

at improving software supply chain governance across organizations of varying maturity, underscoring both the practical value of SBOMs and the challenges of implementing them consistently.[14] These efforts motivate OntoCRA-NS pipelines that treat SBOM availability and structure as compliance-relevant signals, while also highlighting the need for evidence extraction and validation across heterogeneous documentation.

Across these research threads—RegTech compliance automation, ontology/rule-based assessment, cybersecurity certification ontologies, and SBOM standardization—the common gap is a CRA-specific, end-to-end mechanism that connects real technical documentation to explainable compliance outcomes (including conservative handling of incomplete evidence) and downstream artifacts such as draft Declarations of Conformity. The CRA-NS proof-of-concept addresses this gap by integrating extraction, ontology population, and deterministic reasoning into a single auditable pipeline tailored to CRA conformity triage.

RegTech research has demonstrated the potential of ontology-based and rule-driven approaches for automating regulatory compliance checks while preserving auditability and explainability. The CRA-NS framework follows this paradigm by encoding CRA requirements into a formal ontology and applying deterministic reasoning to support conformity assessment decisions.

IV. USE-CASE: PRE-MARKET CONFORMITY DOSSIER TRIAGE

We consider a manufacturer preparing a CRA conformity dossier for a product with digital elements prior to EU market placement. The dossier typically includes technical documentation, user documentation, security policy statements, SBOM (e.g., SPDX/CycloneDX format), vulnerability information (e.g., CVE – Common Vulnerability Exposure references and CISA KEV – Known Exploitable Vulnerability flags), etc. [1] OntoCRA-NS is used as a triage assistant: it ingests the dossier, extracts CRA-relevant attributes into Evidence JSON, populates OntoCRA v3, and produces an explainable outcome indicating whether the dossier appears to satisfy the implemented CRA requirements (*DoC_Granted*), violates one or more requirements (*BannedFromMarket*), or lacks sufficient evidence to reach a decision. [2]

The outcome is accompanied by a justification trace connecting source evidence to rule premises, enabling compliance officers (or developers) to remediate gaps (e.g., *missing SBOM, insufficient support duration, presence of known exploited vulnerability*) before formal submission. This aligns with the framework's intent to assist—not replace—human evaluators while maintaining auditability.

Table 1 PoC decision based on the CRA legal status

PoC Decision Outcome	Ontology-Inferred Decision	Triggering Conditions (PoC Rules)	Interpretation
Compliant	DoC_Granted_Decision	Required conformity evidence present (e.g., SBOM and Technical Documentation) and security update support period meets or exceeds expected product lifetime	Product is admissible for market placement under CRA triage criteria
Non-compliant	BannedFromMarket_Decision	One or more critical conformity requirements violated (e.g., missing SBOM, missing technical documentation, known exploited vulnerability, insufficient security update support)	Product is not admissible for market placement under CRA triage criteria
Assessment Required	*No decision inferred*	Evidence incomplete or insufficient to trigger any decision rule	Case requires human assessment and cannot be automatically triaged

In the PoC, the reasoner may infer an explicit compliance decision (DoC_Granted_Decision) or non-compliance decision (BannedFromMarket_Decision) (Table 1). When the knowledge base lacks sufficient evidence to satisfy either decision rule set, no decision is inferred by the reasoner. We interpret this third outcome as Assessment Pending, reflecting a realistic conformity triage state where additional evidence or clarification is required prior to drafting a defensible Declaration of Conformity.

V. ONTOCRA-NS PROOF OF CONCEPT SYSTEM

System overview is shown in Figure 1 and detailed described in [2].

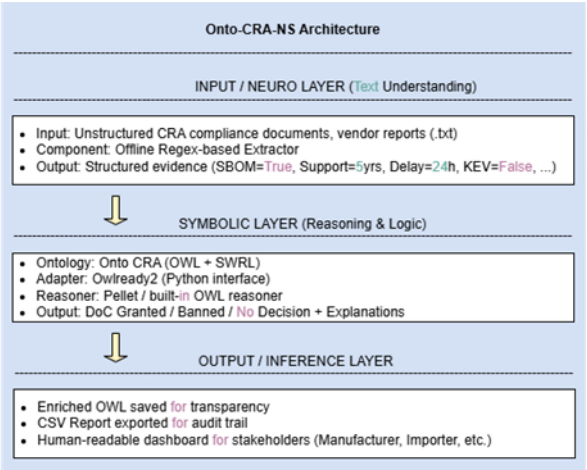

Figure 1 High-level architecture [2]

A. INPUT LAYER: DOSSIER ACQUISITION AND NORMALIZATION

The PoC ingests a product dossier consisting of heterogeneous evidence sources including product documentation/technical documentation, SBOM files, vulnerability reports (CVE lists and KEV indicators, support-policy statements, etc. (From [1]). Evidence may be free-text or semi-structured (PDF, JSON, XML). All textual sources are normalized into plain text for extraction.

B. EVIDENCE JSON SCHEMA (CONTRACT)

In this PoC, the Evidence JSON is instantiated in a simplified, flat attribute-value form sufficient to populate OntoCRA v3. While the OntoCRA-NS architecture supports richer evidence objects (including provenance spans and confidence scores), these extensions are left to future work and LLM-assisted extraction. The current implementation preserves traceability by linking inferred facts back to source documents at report-generation time. When dossier documents provide conflicting statements (e.g., different support periods) or when extracted values cannot be grounded to an explicit text span, the PoC assigns the corresponding evidence attribute as unknown rather than forcing a value. Under open-world semantics this may yield Assessment Pending, which is preferable to unsafe compliance inference in regulatory settings.

C. NEURAL EXTRACTION MODULE

The OntoCRA-NS neural layer acts as the perception component that interprets unstructured dossier documents and extracts CRA-relevant attributes. In the foundational prototype, this layer was simulated using a rule-based extractor that detects key compliance attributes and produces Evidence JSON.

Table 2 Mapping table (paper-ontology-code)

Evidence concept	Evidence JSON field	OntoCRA property	Rule impact
SBOM presence	hasSBOM	*hasSBOM*	Non-compliance if missing
Exploited vuln.	knownExploit	*hasKnownExploit*	Immediate ban
Support duration	supportPeriodYears	security update support period (years)	Ban if < 5
Tech. documentation	hasTechnicalDocumentation	*hasTechnicalDocumentation*	Non-compliance if missing

D. ONTOLOGY ADAPTER (OWLREADY2)

The ontology adapter bridges neural extraction and symbolic reasoning. It parses Evidence JSON and programmatically populates the ontology using Owlready2 by creating or updating individuals of class *ProductWithDigitalElements* with corresponding properties (e.g., *hasSBOM, hasTechnicalDocumentation, supportPeriodYears*, etc). The adapter enforces type consistency and validation, then triggers reasoning via the Pellet API.

Figure 2 illustrates the end-to-end CRA-NS proof-of-concept pipeline, showing how technical documentation is transformed into structured evidence, ontology instances, and compliance decisions

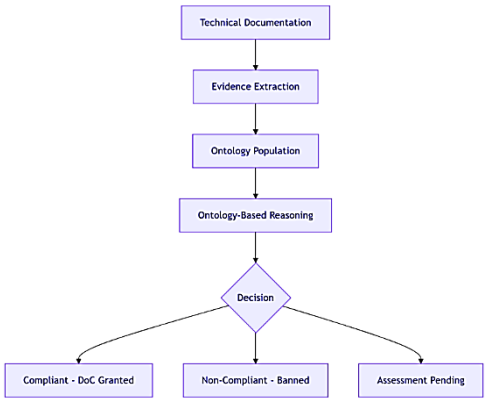

Figure 2 OntoCRA-NS Pipeline

VI. SYMBOLIC REASONING (ONTOCRA V3, SWRL, PELLET)

The reasoning core relies on the OntoCRA v3 ontology formalizing CRA Essential and Vulnerability Requirements in OWL and SWRL, executed by the Pellet reasoner. The ontology encoding the implemented CRA requirements. Reasoning logic includes direct mappings such as "missing SBOM → Non-Compliant," "missing technical documentation → Non-Compliant," and "known exploited vulnerability → Non-Compliant." When all implemented requirements are satisfied and the declared support duration is at least five years, the system classifies the dossier as Compliant (i.e., sufficient basis for issuing a DoC).

Assessment Pending is not explicitly encoded as an OWL class in OntoCRA v3. Instead, it is an operational interpretation applied when the reasoner infers neither a compliance nor a non-compliance decision. This design choice preserves conservative reasoning semantics and avoids introducing default assumptions under incomplete evidence.

In this PoC, Assessment Pending is returned if and only if (i) neither DoC_Granted_Decision nor BannedFromMarket_Decision is entailed by Pellet under OntoCRA v3 and (ii) at least one decisive predicate required by the compliance rule set is unknown or ungrounded in the dossier. This definition intentionally avoids default assumptions and reflects conservative CRA triage practice.

Scope Disclaimer: OntoCRA v3 encodes only a subset of CRA Essential and Vulnerability Requirements sufficient to demonstrate the CRA-NS reasoning pipeline. Several extracted attributes (e.g., secure-by-design statements or risk assessment descriptions) are currently treated as non-decisive contextual evidence. Extending these attributes into decision-triggering rules is left to future ontology revision.

The current rule set focuses on documentary completeness (SBOM, technical documentation, VHP/incident reporting), lifecycle support sufficiency, and severe vulnerability evidence (exploited vulnerability). Controls such as secure update mechanism quality or secure-by-design maturity are extracted as contextual attributes but are not yet decision-triggering in OntoCRA v3. Consequently, dossiers may remain Assessment Pending if they provide partial evidence not captured by decisive rules.

A. OUTPUT GENERATION AND FEEDBACK LOOP

After reasoning, OntoCRA-NS generates a compliance decision (e.g., DoC_Granted or BannedFromMarket) together with an explanation in natural language or query-based form. Explanations link inferred outcomes to rule premises, enabling traceability and auditability. A report is generated in readable .pdf format, and DoC (Declaration of Conformity) for compliant product is generated. A human reviewer can validate or correct outputs; this feedback can be used to improve extraction over time, consistent with the CRA-NS neuro-symbolic learning loop.

VII. POC SETUP

This section describes the experimental setup used to validate the proposed CRA-NS proof of concept (PoC). The PoC integrates symbolic reasoning over regulatory requirements with structured evidence extraction from technical documentation.

The implementation uses OntoCRA v3 from https://github.com/jascosic/OntoCRA-NS as the formal knowledge representation of CRA requirements, with Pellet as the OWL reasoner. Evidence extraction and ontology population are implemented in Python using Owlready2. The system processes product dossiers and produces structured compliance reports, including a machine-readable decision trace and a human-readable PDF explanation of inferred compliance outcomes.

All experiments were executed on a standard notebook environment (Windows 11 Pro, Intel i9 CPU, 32 GB RAM). As runtime performance is not the focus of this PoC, hardware characteristics are reported for reproducibility only.

For the PoC was used:

- Ontology version (OntoCRA v3) – available here: https://github.com/jascosic/OntoCRA-NS
- Reasoner (Pellet)
- Tooling (Owlready2, Python script)
- Dataset size (6 dossiers – technical documentation files from the CURIUM project)
- Outputs produced (report in .pdf format, in cmd + DoC in readable .pdf format)

Table 3 CRA obligation mapping

CRA Obligation (Article / Annex)	Ontology Concept	SWRL Rule
SBOM availability	SBOM, hasSBOM	Rule 1
Technical documentation	TechnicalDocumentation	Rule 2
Known exploited vulnerabilities	Vulnerability, KnownExploitedVulnerability	Rule 3
Security update support	supportPeriodYears, expectedLifetimeYears	Rules 4–7

The mapping from Table 3 reflects a representative subset of CRA requirements and is intended for demonstrative purposes.

A. Dataset and Dossier Composition

We construct a dataset of six synthetic but realistic product dossiers from the CURIUM project, each representing the technical documentation of a distinct digital product subject to the CRA. Each dossier aggregates heterogeneous document types corresponding to CRA-relevant evidence categories, including product documentation with SBOM artifacts, vulnerability disclosures (CVE and KEV indicators), support policy statements, and user documentation.

The dossiers are intentionally designed to trigger different regulatory outcomes (*Compliant, Non-Compliant, Assessment Pending*) and to exercise key OntoCRA v3 rule conditions such as missing SBOMs, presence of known exploited vulnerabilities, and insufficient support duration. Each dossier is processed independently by the CRA-NS pipeline, from evidence extraction to ontology-based reasoning and decision explanation. Table 4 summarizes the expected regulatory scenarios and target outcomes for each dossier.

Table 4 Products with expected scenarios

Dossier (Product)	Expected scenario	System outcome	Reason(s)
CAC	Exploited vulnerabilities present	**Non-Compliant**	Reason_ExploitedVulnerability
CyReA	Full evidence, no exploited vulns	**Compliant**	Reason_Compliant
DPMA	Evidence gaps / update mechanism weakness	**Assessment Pending**	None inferred
DPRA	Full evidence	**Compliant**	Reason_Compliant
PSTV	Full evidence	**Compliant**	Reason_Compliant
DEMF	Partial/unclear compliance claims	**Assessment Pending**	None inferred

B. Ground Truth Labeling Protocol

Ground-truth labels are defined at the level of extracted evidence rather than at the level of regulatory interpretation. For each dossier, human annotators manually identify and record ground-truth values for the Evidence JSON attributes implemented in the PoC, along with explicit references to supporting documents and text excerpts.

Ground-truth compliance outcomes are then deterministically derived by applying the same OntoCRA v3 rule logic used by the reasoning engine to the ground-truth evidence values. This design ensures that evaluation focuses on the correctness of evidence extraction and ontology population, while avoiding subjective interpretation of regulatory compliance. As a result, discrepancies between system outputs and ground truth directly reflect extraction or representation errors rather than semantic ambiguity in the regulatory rules.

Importantly, the system is not evaluated on its ability to "rediscover" the rules; rather, it is evaluated on whether it extracts and grounds the evidence values that serve as rule inputs. This isolates extraction and ontology population errors as the primary sources of end-to-end decision mismatch.

C. Metrics

The PoC is evaluated using task-appropriate, deterministic metrics aligned with regulatory decision support rather than statistical classification performance.

Evidence Extraction Accuracy.

For each dossier, extracted Evidence JSON fields are compared against manually annotated ground-truth values. Evaluation is performed at the field level, and each field is classified as correct or incorrect based on exact match (for boolean and categorical attributes) or value agreement within predefined tolerance (for numeric attributes such as declared support duration). Results are reported per dossier and per evidence category to support transparent error analysis.

End-to-End Compliance Outcome Agreement.

At the reasoning level, the inferred compliance outcome (*Compliant, Non-Compliant, Assessment Pending*) is compared against the ground-truth outcome derived from annotated evidence using the same OntoCRA v3 rules. Outcome agreement is reported as the number of dossiers for which the system inference matches the expected regulatory decision.

False Compliant Cases (Safety Metric).

Because regulatory automation must prioritize the avoidance of unjustified compliance declarations, we explicitly report false compliant cases, defined as instances where the system infers a Compliant outcome while the ground-truth outcome is *Non-Compliant*. This metric reflects a critical safety property of the system and is particularly relevant in the CRA context.

Given the small scale and exploratory nature of the PoC, no aggregate statistical performance measures are reported. Instead, results are presented in a dossier-level breakdown to enable qualitative inspection of system behaviour and reasoning traceability.

D. Baselines and Ablations

To isolate the impact of evidence extraction on symbolic reasoning, all experiments use a fixed ontology (OntoCRA v3) and an identical set of SWRL rules. The baseline configuration employs a deterministic, rule-based evidence extractor implemented in the initial CRA-NS prototype.

No learning-based extractor is evaluated in this work. However, the PoC is explicitly designed to support alternative extractors, including LLM-assisted components, without modifying the symbolic reasoning layer. This modularity aligns with the CRA-NS design principle that neural components may evolve independently while preserving stable, auditable regulatory semantics. An empirical comparison with LLM-assisted extraction is left for future work.

VIII. RESULTS

A. Extractions results

We report field-level extraction outcomes for the six evidence attributes from the project, used by the OntoCRA v3 rules. Because the dataset is intentionally small (six dossiers), results are interpreted case-by-case and aggregated only as descriptive statistics. Table 5 summarizes extraction quality for the baseline rule-based extractor. In addition to correctness of values, we log whether each extracted value is backed by explicit evidence span in the source dossier; ungrounded outputs are treated as unknown to preserve conservative behaviour.

B. End-to-end compliance decision results

Table 5 summarizes end-to-end decision quality. For each dossier, the extracted Evidence JSON is mapped to OntoCRA v3 and reasoned over with Pellet, producing a deterministic outcome. We evaluate whether the inferred outcome matches the ground-truth decision derived from the same rule logic and annotated evidence. We also report a false compliant rate, defined as cases where the system outputs Compliant while ground truth is *Non-Compliant*; this is treated as a critical error.

Table 5 End-to-end compliance decision results

Product	Assurance Level	Decision	Reason(s)	Interpretation
CAC	Basic	Banned from Market	Exploited Vulnerability	Non-Compliant
CyReA	Basic	DoC Granted	Compliant	Compliant
DPMA	Basic	No decision	—	Assessment Pending
DPRA	Basic	DoC Granted	Compliant	Compliant
PSTV	Basic	DoC Granted	Compliant	Compliant
DEMF	Basic	No decision	—	Assessment Pending

While this PoC is evaluated on six dossiers, the runtime breakdown suggests that ontology population and DL

reasoning are lightweight relative to evidence extraction. The expected scalability bottleneck in practical deployments is document ingestion and extraction across large dossier collections, while the symbolic reasoning layer remains comparatively stable as rule coverage expands.

C. Explainability/Traceability results

1) End-to-End Decision Outcomes

The reasoner inferred three compliant decisions (DoC_Granted_Decision), one explicit non-compliant decision (BannedFromMarket_Decision), and two cases where no decision was inferred. The latter are interpreted as Assessment Pending, indicating that the available evidence was insufficient to satisfy either compliance or non-compliance rules under OntoCRA v3.

2) Non-Compliance Detection (CAC)

The CAC dossier was classified as Non-Compliant due to the presence of at least one known exploited vulnerability. The ontology-based reasoning process inferred BannedFromMarket_Decision with Reason_ExploitedVulnerability, demonstrating the system's ability to elevate high-severity security evidence into deterministic regulatory outcomes consistent with CRA safety objectives.

3) Compliant Dossiers (CyReA, DPRA, PSTV)

CyReA, DPRA, and PSTV were classified as compliant and triggered DoC_Granted_Decision. All three dossiers contained evidence for SBOM presence, technical documentation completeness, vulnerability handling and incident reporting processes, and sufficient security update support relative to expected product lifetime as modeled by OntoCRA v3. The system automatically generated draft DoC artifacts for these products, demonstrating end-to-end operational viability.

4) Assessment Pending (DPMA, DEMF)

Cases where no decision is inferred by the rule base are interpreted as requiring human assessment and are reported as 'undecided' without constituting a formal decision class.For DPMA and DEMF, the reasoner inferred no final decision. This outcome reflects Assessment Pending, indicating that while some CRA-relevant evidence was present, the available facts were insufficient to satisfy the decisive rule conditions encoded in OntoCRA v3. Importantly, the system does not infer compliance by default, thereby avoiding unsafe or unjustified regulatory conclusions. (Table 6)

Table 6 Missing or non-decisive evidence leading to assessment pending

Product	Blocking factor (from ontology perspective)
DPMA	Secure update mechanism and/or lifecycle support not encoded as decisive rules
DEMF	Ambiguous compliance claims; insufficient structured evidence for decisive predicates

D. Runtime Performance

Runtime performance was on the way that measure runtime per dossier for extraction, ontology population, and reasoning. Table 7 reports average timings. While the dataset is small, this breakdown isolates bottlenecks and confirms that symbolic reasoning remains lightweight compared to document extraction. Prior work reported sub-2-second reasoning time per product; this study reports updated timings under dossier-based ingestion.

All runtime measurements are reported as average per-dossier values. Pellet reasoning is executed once per batch and amortized over the number of processed dossiers.

Table 7 Runtime per dossier

Stage	Avg time per dossier (s)	Notes
Extraction	< 0.001	Rule-based baseline extractor
Ontology population	0.001	Owlready2 individual creation
Reasoning	0.214	Pellet; batch reasoning amortized
Total	0.215	End-to-end CRA-NS pipeline

IX. Case-Study analysis and interpretation

These case studies emphasize traceability from dossier artifacts to extracted Evidence JSON, ontology facts, and rule-triggered outcomes.

A. DOSSIER CYREA: COMPLIANT DOSSIER

The CyReA dossier represents a product with digital elements providing tooling for cybersecurity and resilience assessment. The technical documentation includes a structured SBOM, comprehensive technical documentation, and explicit descriptions of vulnerability handling and incident reporting processes. The declared post-market security update support period is consistent with the expected product lifetime, satisfying the lifecycle support requirements encoded in OntoCRA v3. During evidence extraction, the CRA-NS pipeline identified all decisive compliance predicates required by the modeled CRA rules. No known exploited vulnerabilities were reported in the dossier, and no vulnerability evidence triggered non-compliance conditions. The populated ontology instance therefore satisfied the compliance rule set, leading the reasoner to infer a *DoC_Granted_Decision* with the associated justification *Reason_Compliant.* As part of the PoC workflow, a draft Declaration of Conformity (DoC) document was automatically generated for CyReA. This case demonstrates the ability of the CRA-NS framework to support end-to-end conformity assessment, from unstructured technical documentation to a defensible compliance decision and downstream regulatory artifact generation. Importantly, the decision remains explainable and traceable, with each inferred conclusion grounded in explicitly modelled evidence rather than heuristic scoring.

B. DOSSIER CAC: NON-COMPLIANT

Although the CAC dossier includes an SBOM, technical documentation, and a declared five-year support period, it explicitly reports the presence of multiple vulnerabilities under active exploitation. Under the modelled CRA logic, such evidence constitutes a decisive non-compliance condition, leading to a *BannedFromMarket_Decision* irrespective of other documentary completeness indicators.

C. *DOSSIER DPMA: ASSESSMENT PENDING DUE TO INCOMPLETE EVIDENCE*

The DPMA dossier illustrates a realistic borderline case in which partial CRA-relevant evidence is present, but not in a form sufficient to trigger a definitive compliance or non-compliance decision under the current ontology and rule coverage. The technical documentation includes descriptive material regarding the product's purpose and maturity assessment capabilities; however, several decisive predicates required by OntoCRA v3 could not be conclusively populated during evidence extraction. Specifically, while DPMA contains general compliance-oriented statements, the extracted evidence did not satisfy the complete set of conditions required for inferring DoC_Granted_Decision. At the same time, no decisive non-compliance condition—such as the presence of known exploited vulnerabilities or an explicitly insufficient support period—was detected. As a result, the reasoner inferred no compliance decision. This outcome is interpreted as Assessment Pending, reflecting a conservative reasoning stance aligned with OWL's open-world semantics. In Table 8 are shown all Dossier scenarios with expected CRA-NS outcomes.

Table 8 (All products with expected outcomes)

Dossier ID	Product	Scenario design	Key missing / triggered evidence	Expected outcome
D1	CAC	Explicit security risk scenario	Presence of known exploited vulnerabilities (KEV); CVEs under active exploitation	Non-Compliant (Banned from Market)
D2	CyReA	Fully documented compliant product	SBOM present; complete technical documentation; vulnerability handling process; sufficient security update support; no known exploited vulnerabilities	Compliant (DoC Granted)
D3	DPMA	Partial / ambiguous documentation	Insufficient decisive evidence to satisfy compliance rules; no triggered non-compliance conditions	Assessment Pending
D4	DPRA	Fully documented compliant product	SBOM present; technical documentation; vulnerability handling; lifecycle support aligned with expected lifetime	Compliant (DoC Granted)
D5	PSTV	Fully documented compliant product	SBOM present; technical documentation; vulnerability handling; sufficient support period; no exploited vulnerabilities	Compliant (DoC Granted)
D6	DEMF	Incomplete or weakly structured compliance evidence	Missing or non-decisive predicates for compliance; no explicit non-compliance trigger	Assessment Pending

X. DISCUSSION AND THREATS TO VALIDITY (LIMITATION)

OntoCRA v3 encodes a subset of CRA Essential and Vulnerability Requirements sufficient to demonstrate end-to-end neuro-symbolic reasoning. Attributes such as secure-by-design maturity or secure update mechanisms are extracted as contextual evidence but are not yet decisive in the current rule set. Extending these attributes into additional compliance rules is planned future work. Viewed as a RegTech system, CRA-NS does not aim to replace notified bodies or human assessors, but to support them by providing structured, explainable, and scalable compliance triage under the Cyber Resilience Act. Dataset size and representativeness: The study uses six intentionally designed dossiers to cover diverse outcomes. This supports feasibility and traceability analysis but does not provide statistical generalization.

The dossier-based PoC demonstrates how CRA-NS can support pre-market conformity triage by transforming heterogeneous documentation into structured evidence and applying deterministic reasoning to reach auditable outcomes. A key practical benefit is conservative handling of uncertainty: when required evidence is missing or cannot be grounded to source text, the system returns Assessment Pending rather than granting compliance. This aligns with compliance practice, where additional documentation is requested before a Declaration of Conformity is drafted or accepted.

Some of the limitations and open questions for future research are:

- Extraction uncertainty: extractors may fail on nuanced or implicit statements (e.g., vague support commitments). Evidence-grounding checks reduce the risk of unsupported values but may increase Assessment Pending outcomes.
- Partial CRA coverage: OntoCRA v3 encodes a subset of CRA Essential and Vulnerability Requirements. Outcomes therefore reflect compliance with the implemented rules rather than full CRA conformity.
- Rule maintenance: Regulatory interpretation and guidance may evolve. Maintaining alignment requires governance of the SWRL rule set and versioning of ontology releases, particularly when extending beyond the initial requirement subset.

XI. CONCLUSION AND FUTURE WORK

This paper presented an ontology-based proof-of-concept system for pre-market conformity triage under the

Cyber Resilience Act, demonstrating an executable instantiation of the OntoCRA-NS architecture. By integrating structured evidence extraction, ontological modelling, and rule-based reasoning, the proposed approach enables transparent and explainable assessment of selected CRA conformity conditions. The results confirm that symbolic reasoning can effectively support regulatory triage by identifying compliant and non-compliant cases while preserving human oversight for situations requiring further assessment. The work intentionally focuses on a limited subset of CRA obligations and employs deterministic extraction methods in order to validate the consistency and correctness of the reasoning pipeline. As such, the system does not aim to automate legal compliance decisions but to provide structured decision support for regulators and economic operators. Future work will extend the ontology to cover additional CRA requirements, formalize intermediate decision states, and integrate neural or large-language-model-based extraction components to enhance scalability and robustness. Overall, the presented proof-of-concept illustrates the potential of ontology-driven, explainable AI systems as a foundation for trustworthy regulatory technologies in the context of emerging cybersecurity regulation.

ACKNOWLEDGMENT

The research leading to these results received funding from the European Union's HORIZON Innovation Action program GA No. 101120684 - project CUSTODES, as well as DIGITAL-ECCC-2024-DEPLOY-CYBER-06 under proposal number 101190372, project CURIUM. Jasmin Cosic from DEKRA SE (Germany) participates as a CUSTODES project partner and Miroslav Baca from Cyber-security, Ltd. (Croatia) as CURIUM project partner.

We want to thank to all project partners for constructive suggestions on Framework and valuable comments.

Early linguistic polishing was performed with ChatGPT (OpenAI), and final grammar and style adjustments with Microsoft Copilot (M365 Copilot). AI tools were used exclusively for language editing; no technical content, modelling, analyses, or conclusions were produced by AI. All scientific content originates from the authors and the project real pilots/use-cases and was fully reviewed and approved by the authors.

REFERENCES

[1] European Commission, "Cyber Resilience Act," *Digital Strategy*, Dec. 03, 2025. [Online]. Available: https://digital-strategy.ec.europa.eu/en/policies/cra-summary (accessed Feb. 10, 2026)

[2] J. Ćosić, A. Jukan, and M. Bača, "OntoCRA-NS: A Neuro-Symbolic Framework for Explainable and Auditable Cybersecurity Compliance Aligned with EU Regulatory Frameworks," in Proc. 2026 Int'l Conf. on Artificial Intelligence (CAI), Granada, Spain, 2026

[3] D. W. Arner, J. N. Barberis, and R. P. Buckley, "FinTech, RegTech, and the reconceptualization of financial regulation," Northwestern Journal of International Law & Business, vol. 37, no. 3, pp. 371–413, 2017.

[4] Butler, T., O'Brien, L. (2019). Understanding RegTech for Digital Regulatory Compliance. In: Lynn, T., Mooney, J., Rosati, P., Cummins, M. (eds) Disrupting Finance. Palgrave Studies in Digital Business & Enabling Technologies. Palgrave Pivot, Cham. https://doi.org/10.1007/978-3-030-02330-0_6

[5] Lallas, E.N.; Santouridis, I.; Mountzouris, G.; Gerogiannis, V.C.; Karageorgos, A. An SQWRL-Based Method for Assessing Regulatory Compliance in the Pharmaceutical Industry. Appl. Sci. 2022, 12, 10923. https://doi.org/10.3390/app122110923

[6] Horrocks, I., Patel-Schneider, P. F., Polleres, A., Axelrod, V., Berners-Lee, T., Chaudhri, V., ... & Zimmermann, M. (2004). SWRL: A semantic web rule language combining OWL and RuleML (W3C Submission). The World Wide Web Consortium (W3C). www.w3.org

[7] O'Connor, M., & Das, A. (2009). SQWRL: A query language for OWL. In Proceedings of the 6th International Conference on OWL: Experiences and Directions - Volume 529 (pp. 208-215). CEUR-WS.org. https://ceur-ws.org/Vol-529/owled2009_submission_42.pdf

[8] Lallas, E.N.; Santouridis, I.; Mountzouris, G.; Gerogiannis, V.C.; Karageorgos, A. An SQWRL-Based Method for Assessing Regulatory Compliance in the Pharmaceutical Industry. Appl. Sci. 2022, 12, 10923. https://doi.org/10.3390/app122110923

[9] Echhofer, S., van Harmelen, F., Hendler, J., Horrocks, I., McGuinness, D., Patel-Schneider, P., & Stein, L. A. (2004). OWL Web Ontology Language Reference. World Wide Web Consortium (W3C). https://www.w3.org/TR/owl-ref

[10] Butler, T., Brien, L., Understanding RegTech for Digital Regulatory Compliance, Disrupting Finance2019

[11] Cosic, J. Deciphering Cyber-Security Certifications: An Ontological Journey through Composite Systems and their certification, 2024

[12] Bruno Lourenco, et al., Structuring Security: A Survey of Cybersecurity Ontologies, Semantic Log Processing, and LLMs Application

[13] NTIA, "The Minimum Elements For a Software Bill of Materials (SBOM)," National Telecommunications and Information Administration, U.S. Department of Commerce, Jul. 12, 2021. [Online]. Available: https://www.ntia.gov/files/ntia/publications/sbom_minimum_elements_report.pdf

[14] European Union Agency for Cybersecurity (ENISA), "SBOM Landscape Analysis: Towards an Implementation Guide," Public Draft v1.2, Dec. 2025. [Online]. Available: www.enisa.europa.eu. Accessed: Feb. 10, 2026.

[15] CURIUM EU funded project, https://curium-project.eu/ , accessed 26.12.2025

[16] CUSTODES EU funded project, https://custodes-project.eu/, accessed 20.12.2025