Proceedings of the Seventh International Conference on
Railway Technology: Research, Development and Maintenance
Edited by: J. Pombo
Civil-Comp Conferences, Volume 15, Paper 20.3
Civil-Comp Press, Edinburgh, United Kingdom, 2026
ISSN: 2753-3239, doi: 10.4203/ccc.15.20.3

Ontology-Driven Scoping and Prioritisation for Railway Penetration Testing

M. Shapoval[1,2], V. Deniau[1], C. Gransart[1] and A. Lang[2]

[1] Gustave Eiffel University, COSYS-LEOST, Villeneuve-d'Ascq, France
[2] IKOS, Paris, France

Abstract

This paper presents an OWL/SWRL-based ontological framework for reproducible, architecture-driven scoping and prioritisation of penetration testing in railway signalling systems. The framework integrates concepts from IEC 62443, EN 50159, EN 50126 and CLC/TS 50701 into a domain-level terminological box (TBox), and operationalises them through a four-layer Semantic Web Rule Language (SWRL) rule architecture that systematically infers vulnerable communication flows, Safety Integrity Level (SIL)-critical attack patterns and test-coverage metrics from instantiated signalling architectures. A demonstration on a European Train Control System (ETCS) Level 2 architecture shows that the reasoner derives meaningful, standards-justified penetration-testing priorities that are not obtainable from existing standards-agnostic approaches. The framework is intended as a formal assessment layer supporting, rather than replacing, expert judgement in railway signalling security testing.

Keywords: railway signalling, cybersecurity, penetration testing, ontology, industrial control systems, safety integrity level

1 Introduction

Railway signalling and train control systems have experienced a rapid digital transformation, moving from largely isolated, hardware-centric architectures to IP-connected, software-intensive systems increasingly integrated with corporate IT networks and cloud-based platforms. This evolution has expanded the cyber attack surface of railway signalling systems while preserving their safety-critical nature. Successful cyber attacks can disrupt traffic management, degrade signalling integrity, or in the worst case undermine safety functions designed to prevent collisions and overspeed incidents [1, 2]. Recent surveys of railway cybersecurity practice highlight that security assurance is still dominated by documentation-driven risk assessments and high-level compliance with generic standards, rather than systematic, architecture-aware security testing [2, 3].

The regulatory landscape is converging on explicit requirements for structured security testing, including penetration testing, in both industrial and railway contexts. IEC 62443-4-1 – the secure product development lifecycle standard for industrial automation and control systems (ICS) – mandates penetration testing as part of security verification and validation [4]. CLC/TS 50701 – the railway-specific cybersecurity specification – explicitly requires penetration testing to identify vulnerabilities that an attacker could exploit [5]. The NIS2 Directive further requires essential entities, including rail transport operators, to assess the effectiveness of their cybersecurity risk-management measures [6].

However, none of these standards prescribes how penetration testing should be scoped for concrete railway signalling architectures. Railway signalling is additionally constrained by EN 50159 safety-related communication requirements and by SIL classifications, which limit how aggressively live systems can be tested. Any scoping mechanism must therefore account for both security threats and safety constraints simultaneously. In practice, this creates a substantial gap: standards require penetration testing of complex, safety-constrained systems, but provide no reproducible, architecture-driven mechanism to determine what to test and in what order [2, 8].

Existing approaches – including ICS risk assessment standards, adversary technique catalogues, generic penetration-testing methodologies, and ontology-based security frameworks – address parts of this problem but none jointly provides railway-specific semantics, architecture-driven attack-path derivation, and formal reproducibility [9–12]. A systematic comparison is provided in Section 2.

This paper addresses the above gap by proposing an ontology-based framework that automatically derives a prioritised, standards-justified set of penetration-testing targets from a formal model of a railway signalling architecture. Because the derivation logic is encoded as rules over a shared ontology, two assessors given the same architecture model obtain the same findings – making the scoping rationale reproducible and auditable [13, 14]. The framework is demonstrated on a representative ETCS Level 2 architecture using publicly documented ERTMS/ETCS reference material [15].

2 Related Work

Table 1 assesses representative existing approaches against four criteria (**Yes/Partial/No**): railway-specific semantics, architecture-driven attack-path derivation, EN 50159/SIL safety-aware reasoning, and formal reproducibility.

IEC 62443-3-2 introduces the zone-and-conduit model, which has become the de facto standard for structuring ICS architectures, while CLC/TS 50701 explicitly applies this model to the railway sector [5, 9]. However, the standard provides no machine-interpretable mechanism that, given a concrete architecture, deterministically enumerates attack paths with explicit preconditions and safety implications. As a result, two analysts applying IEC 62443-3-2 to the same system may produce different scenario sets [8].

MITRE ATT&CK for ICS provides a structured vocabulary for adversary tactics and techniques in industrial environments [10]. Nevertheless, ATT&CK for ICS does not prescribe how to map its techniques onto a particular system architecture, and does not incorporate railway-specific concepts such as EN 50159 transmission categories, SIL-classified safety functions, or TS 50701 threat actors [5, 7].

Generic penetration testing methodologies, such as PTES, OSSTMM, and NIST SP 800-115, provide a framework for planning, executing, and reporting on security tests, primarily in IT contexts [11, 19, 20]. Applying these methodologies to safety-critical operational technology (OT) without domain-specific adaptation poses operational risk, because test depth and aggressiveness are not bounded by SIL classifications or fail-safe dependencies [16, 17].

The closest existing work to our objective is Staves et al. (2024), who propose a risk-based safety scoping approach for adversary-centric security testing on OT systems [16]. They demonstrate that active OT testing is feasible when test aggressiveness is bounded by safety classifications – a principle our framework operationalises formally for railway signalling via SWRL rules and Neo4j-based attack-path analysis. However, their approach does not encode railway domain semantics, EN 50159 communication categories, SIL-classified attack impact, or TS 50701 threat actors. Approaches using ontologies have shown that formal reasoning can be used to make security assessments more systematic [12, 21]. Fenz and Neubauer formalise the ISO 27002 control process in an OWL ontology with SWRL rules for automated compliance determination [12]; Maunero et al. propose an asset-oriented ontology with inference rules for automated risk assessment [21]. These works are methodologically closest to ours; we adopt their principle of ontology-driven, rule-based security assessment and extend it with railway domain semantics, EN 50159 communication categories, SIL-classified safety functions, and automatic attack-path derivation – none of which are present in existing ICS or IoT ontologies.

The analysis shows that no existing approach jointly provides railway-specific semantics, architecture-driven attack-path derivation, and formal reproducibility.

Table 1: Existing approaches vs. railway-specific, architecture-driven attack-path derivation.

Approach	Railway semantics	Arch.-driven paths	EN 50159/SIL aware	Formal reprod.
IEC 62443-3-2 [9]	**No**: generic ICS zones; no railway assets	**Partial**: zone/conduit defined; scenarios manual	**No**: EN 50159 absent; SIL not linked	**No**: two analysts may differ [8]
CLC / TS 50701 [5]	**Partial**: railway assets and actors defined	**Partial**: lifecycle outlined; no formal mechanism	**Partial**: railway context; EN 50159 not integrated	**No**: practitioner discretion
MITRE ATT&CK ICS [10]	**No**: ICS techniques only	**No**: not architecture-bound	**No**: no EN 50159/SIL	**No**: expert judgement
Generic PT [11, 19, 20]	**No**: IT-centric	**No**: expert scope only	**No**: SIL/EN 50159 absent	**No**: process guidance only
OT safety scoping [16]	**No**: generic OT	**Partial**: safety-bounded scope	**Partial**: EN 50159 absent	**Partial**: partially bounded
Attack trees [33, 34]	**No**: domain-agnostic	**Yes**: formal paths	**No**: no EN 50159/SIL	**Partial**: manual model
Control ontologies [12, 36]	**No**: ISO 27002 / IEC 62443	**No**: no attack paths	**No**: no EN 50159/SIL	**Partial**: consistent outputs
Risk ontologies [21, 35, 37]	**No**: generic ICS/IoT	**Partial**: generic risk indicators	**No**: no EN 50159/SIL	**Partial**: consistent outputs
Our proposition	**Yes**: railway zones, EN 50159 flows, SIL functions, TS 50701 actors	**Yes**: auto-inferred VulnerableFlow, Category2/3Flow, SILCriticalPath	**Yes**: EN 50159 categories, SIL, fail-safe, TS 50701 actors	**Yes**: DL-safe SWRL; same ABox ⇒ identical inference

[a] Formal reproducibility: two independent assessors with the same ABox and rule set obtain identical findings.

Table 1 and the analysis above show a two-tier gap. At the methodological level, no existing penetration testing methodology takes into account the needs of railway systems, the EN 50159 communication category, or safety constraints classified by SIL; PTES provides the most adaptable process structure but remains architecture- and domain-agnostic, leaving practitioners without a railway-specific scoping mechanism [1, 2, 11]. At the assessment level, no existing standard, ontology or tool provides a railway-specific, machine-interpretable mechanism that systematically derives attack paths with explicit preconditions and safety implications from a concrete architectural model; and no ontology-based framework has been defined for railway

signalling or linked to penetration testing scope selection [5, 12, 21]. To the authors' knowledge, there is currently no peer-reviewed work that jointly addresses both levels. The framework proposed in this paper addresses the gap at the assessment level by providing a formal, standards-compliant mechanism for scope definition, which is currently lacking in PTES-based railway testing.

3 Ontological Framework

The proposed ontological framework models ICS cybersecurity at the domain level by capturing structural, safety and threat-related concepts in an OWL TBox that is independent of any specific technology family [12, 28]. The TBox vocabulary is organised around three thematic groups: zone/conduit structure grounded in IEC 62443, safety communication properties generalised from IEC 61508 and IEC 62443-3-3, and threat-actor classification derived from MITRE ATT&CK for ICS – all encoded within a single domain-agnostic ICS schema [5, 9, 10, 31]. Railway-specific standards (EN 50159, EN 50126, CLC/TS 50701) informed the design where they provided more precise documentation of concepts common to safety-critical ICS broadly; the resulting vocabulary is not railway-specific and can be instantiated for other OT domains [5, 7, 30].

Figure 1 illustrates the overall ontology structure and workflow: the TBox defines the shared vocabulary of railway signalling concepts; the assertional box (ABox) instantiates it with a concrete signalling architecture (e.g. ETCS). SWRL rules and the OWL DL reasoner operate jointly over both to produce inferred security findings used for penetration-testing scope definition. The TBox is imported into the ABox via `owl:imports`, ensuring that all class and property definitions are available to the reasoner without duplication.

The TBox was initially authored in Protégé 4 for class hierarchy editing and consistency checking; the ABox and SWRL rules are generated programmatically via Python and OWLready2 [40], with HermiT as the underlying OWL DL reasoner.

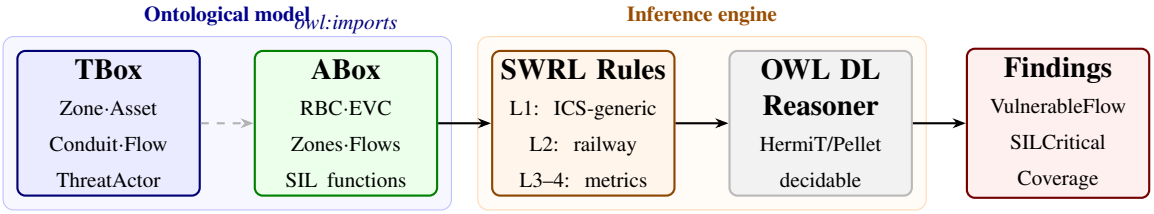

Figure 1: Ontological framework architecture.

3.1 TBox: Three Thematic Groups

The TBox (Figure 2) organises classes into three groups. The **structural group** (IEC 62443-3-2 [9, 29]) provides `Zone`, `Conduit`, `TrustBoundary`, `Asset`, `InformationFlow`, and `Function` to model ICS topology; keeping these generic enables reuse across OT domains. The **safety**

communication group (IEC 61508, IEC 62443-3-3 [29, 31]) encodes `SafetyRelatedFunction` with `safetyIntegrityLevel` (SIL 0–4) and `SafetyRelatedInformationFlow` with defence properties traceable to IEC 62443-3-3 FRs and EN 50159 threats (Table 2). The EN 50159 transmission exposure classification (closed/partly-open/open) enables inferences such as `VulnerableToMasquerade` [7]. The **threat-actor group** (MITRE ATT&CK for ICS, NIST SP 800-82 [10,17]) encodes `ThreatActor` capability tiers (ScriptKiddie through NationState, InsiderOperator, SupplierRepresentative), `AttackVector` means, and three relationships (`hasAccessToZone`, `hasAccessToAsset`, `canInitiateFlow`) for attack-path reasoning.

Table 2: SafetyRelatedInformationFlow defence properties and normative sources.

TBox property	**Source**	**EN 50159 threat**
hasAuthentication	FR3, EN 50159	Masquerade
hasIntegrityProtection	FR3	Corruption
hasSequenceProtection	FR3, EN 50159	Repetition, Resequencing
hasTimeoutMechanism	FR3, EN 50159	Deletion, Delay
hasSafetyCode	EN 50159	Corruption
hasSourceDestinationID	EN 50159	Insertion, Masquerade
hasRedundancy	FR7	Deletion
hasEncryption	FR4	Eavesdropping

FR3=System Integrity; FR4=Confidentiality; FR7=Availability [29].

Figure 2 shows the TBox class diagram; solid arrows indicate associations, open triangles subclass relationships, and dashed arrows structural links.

The TBox is deliberately domain-agnostic – concrete systems appear only as ABox individuals – making the schema reusable across OT domains by substituting only the ABox [9,17,31]. The following section describes how SWRL rules operationalise this TBox into automated security assessment.

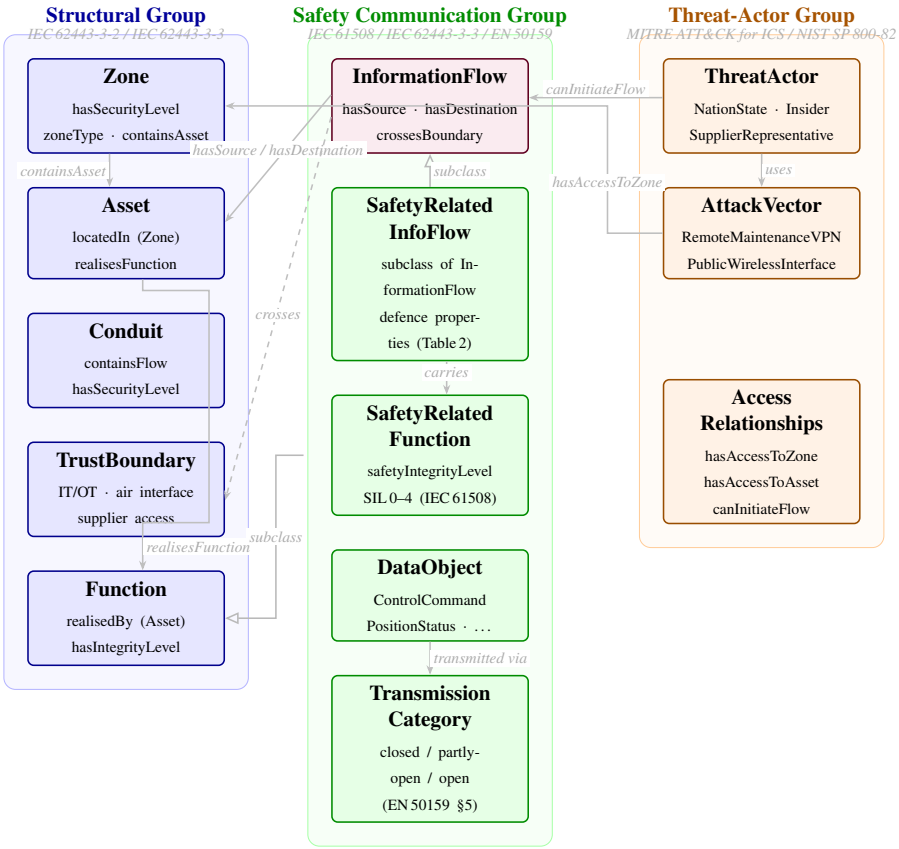

Figure 2: TBox class diagram. InformationFlow (purple, centre) connects all three thematic groups.

3.2 Rule Architecture

To operationalise the TBox into a security assessment process, we encode the inference logic as DL-safe SWRL rules over the shared ontology [13, 14]. The resulting 66 rules form a four-layer monotonic architecture (Figure 3) in which each layer adds inferred facts without retracting earlier ones – organised by a volatility gradient reflecting each normative source's expected rate of change.

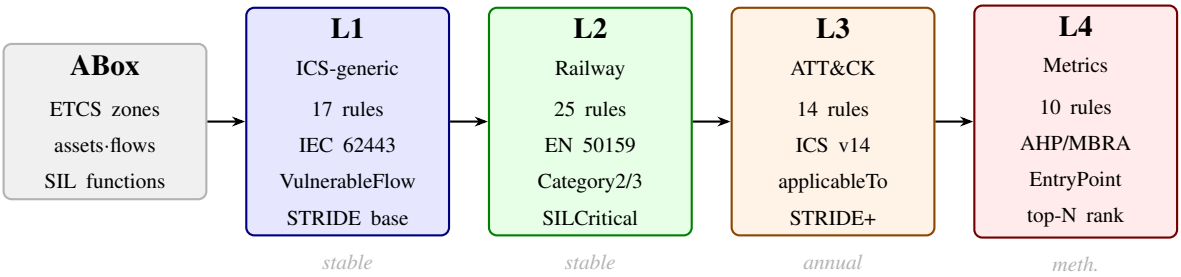

Figure 3: Four-layer SWRL rule architecture (L = Layer).

Layer 1 (17 rules, IEC 62443 FR1–FR7 [29]) flags unprotected cross-boundary

flows, security-level gaps, and external-zone reachability; it also performs first-pass STRIDE classification from TBox properties alone, providing a structural baseline reusable across OT sectors [17].

Layer 2 (25 rules, EN 50159, IEC 61508, TS 50701 [5, 7, 31]) specialises Layer 1 for railway: it classifies flows into EN 50159 Category 2/3; operationalises all seven EN 50159 §5 threat types; elevates findings to safety-critical level when flows carry safety-relevant payloads; derives `FailSafeVulnerability` when the EN 50126 fail-safe barrier is compromised [30, 31]; classifies assets by SIL tier; and flags `HighRiskMaintenancePath` for supply-chain vectors per TS 50701 §7.2.2 [5].

Layer 3 (14 rules) maps vulnerability classes to ATT&CK for ICS v14 [10] via data–logic separation (`attack_techniques.json`); STRIDE enrichment completes a dual-path derivation that elevates confidence when structural and technique-based paths agree [35].

Layer 4 (10 rules) implements an AHP-weighted MBRA pipeline (CR = 0.055 [9, 32]): EntryPoint identification, transitive reachability propagation, SIL-classified impact ranking, and MaximumRiskAsset override when FailSafeVulnerability co-occurs with SIL 3–4 [30, 31]. Neo4j Cypher queries aggregate scores and rank top-N paths.

4 Demonstration on an ETCS Architecture

4.1 ABox Instantiation

The TBox was instantiated for a representative ETCS Level 2 architecture [2, 15, 38] via `generate_abox.py`, producing 424 individuals (assets, zones, flows, safety functions, threat actors); reasoning expands this to 517. Figure 4 shows asset RIU (SIL-2) in the Distributed Trackside Zone carrying inferred classes `HighImpactTarget` and `ReachableAsset` – evidence that the pipeline correctly propagates vulnerability classifications.

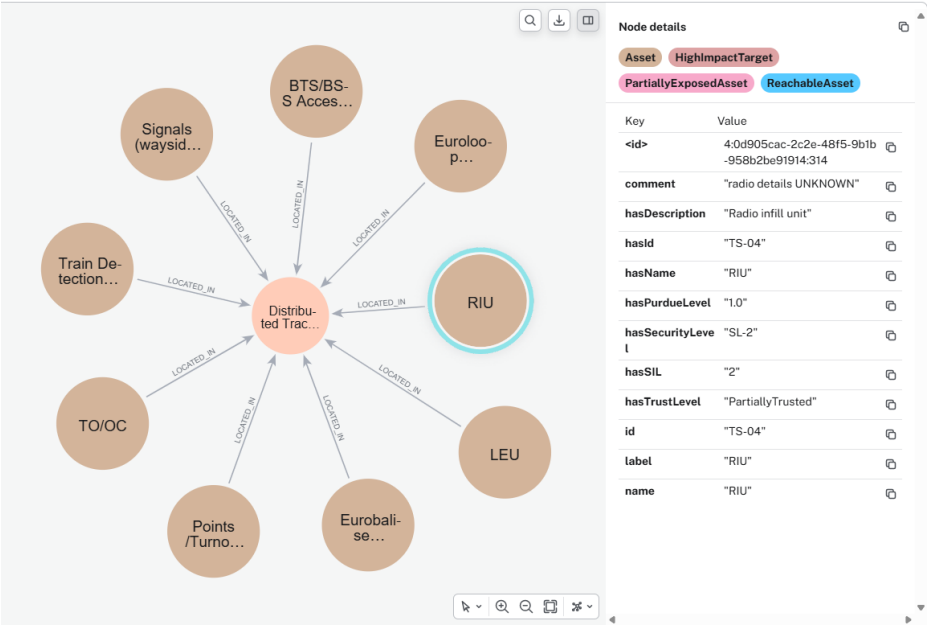

Figure 4: Distributed Trackside Zone – post-reasoning ABox state in Neo4j.

4.2 Inferred Vulnerabilities and Prioritised Flows

The 66-rule pipeline produced 2 120 new triples. Tables 3 and 4 summarise the results: flows classified as EN 50159 Category 2/3 leading to SIL 3–4 assets rank HIGH; non-safety flows rank MEDIUM or LOW.

Table 3: Inference statistics: key inferred classes from the ETCS Level 2 ABox.

Category	Class	Count
Vulnerable flows	VulnerableFlow	58
	Category 2 Flow	50
	Category 3 Flow	7
EN 50159 threats	CorruptionVulnerability	85
	DelayVulnerability	78
	ResequencingVulnerability	82
Safety-critical	MaximumSILRisk	1
	FailSafeVulnerability	1
	CriticalImpactTarget	5
	HighImpactTarget	18
Reachability	EntryPoint	25
	ReachableAsset	28
Maintenance risk	HighRiskMaintenancePath	5

Table 4: Example inferred vulnerabilities and pentest priorities (subset).

Flow	Key inferred classes	L	P
RBC_EVC	Category3Flow, CorruptionVuln., DelayVuln.	2	H
MAINT_VPN	HighRiskMaintPath, PrivAccessRisk, RemoteAccessRisk	2	H
SIL4_Ctx	MaxSILRisk, CritIntegrityViol., FailSafeVuln.	2/4	H

H = HIGH, M = MEDIUM, L = layer, P = priority

4.3 End-to-End Example: EuroRadio

`Flow_RBC_EVC` (EuroRadio, RBC→EVC) is modelled as SafetyRelatedInformationFlow with SIL 4, crossing an AirInterfaceBoundary with authentication and encryption but without integrity or sequence protection [38]. Layer 1 infers `VulnerableFlow`; Layer 2 classifies it as `Category3Flow` + `CorruptionVulnerability` per EN 50159 [7] and derives `MaximumSILRisk` to `CriticalImpactTarget` assets (Figure 5). This gives the penetration tester a reproducible, standards-justified rationale for prioritising integrity and replay-protection testing.

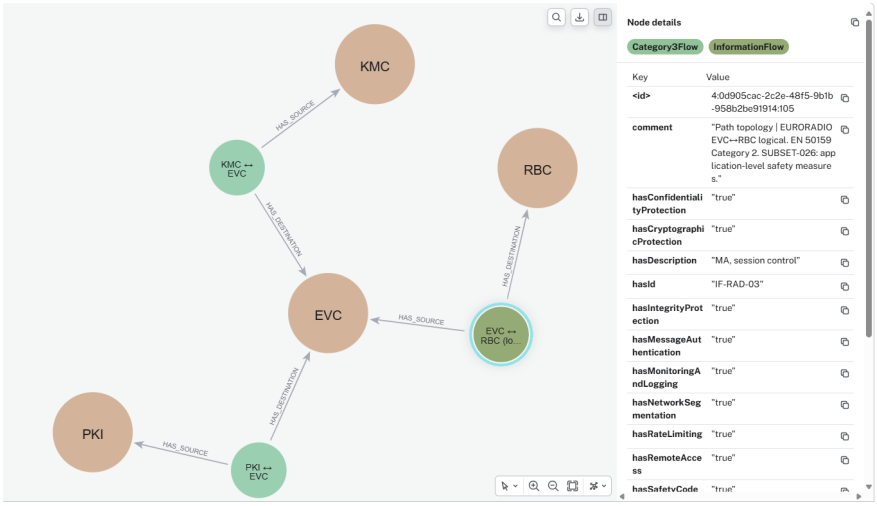

Figure 5: EVC↔RBC EuroRadio flow (IF-RAD-03) in Neo4j.

4.4 Coverage and Prioritisation Metrics

Layer 4 rules and the accompanying Neo4j queries compute summary coverage metrics that support reporting and test planning. Table 5 presents the coverage picture for the ETCS Level 2 ABox.

Table 5: Coverage and risk summary – ETCS Level 2 ABox.

Metric	N	Metric	N
VulnerableFlow (total)	58	EntryPoint	25
EN 50159 Category 2 flows	50	ReachableAsset	28
EN 50159 Category 3 flows	7	HighRiskMaintPath	5
CriticalImpactTarget assets	5	RemoteAccessRisk	14
HighImpactTarget assets	18	MaximumSILRisk	1
FailSafeVulnerability	1	PrivilegedAccessRisk	1

Once TestCase individuals are added, Layer 4 computes Attack Surface Coverage, Threat Coverage (EN 50159 Category 2/3 flows), and SIL-Critical Coverage – a machine-interpretable reporting capability unavailable in standards-agnostic tooling.

5 Discussion

The ontological framework serves a dual purpose: its immediate role is methodological – providing a standards-grounded basis for identifying what PTES requires to adapt to railway signalling contexts – and its longer-term role is operational, serving as a core scoping mechanism within a complete penetration-testing lifecycle [1, 16].

5.1 Limitations

This work has several limitations that delineate the scope of the present contribution and motivate future extensions.

The demonstration is based on a single, albeit representative, ETCS Level 2 railway architecture. Although the TBox is designed at the railway signalling domain level and is intended to be reusable, the current evaluation does not include other families such as Communications-Based Train Control (CBTC), legacy relay-based interlocking, or hybrid systems, and therefore cannot claim empirical generality across all railway contexts [1, 2].

The framework has been validated in terms of internal consistency – reasoner behaviour, competency question coverage, and rule-firing statistics – but has not been evaluated against a ground truth of known, fully disclosed vulnerabilities and incidents. Such validation would require access to sensitive operational data that is typically unavailable in the public domain, and constitutes a recognised challenge for OT security research more broadly [8, 16].

The framework focuses on structural and configuration-level vulnerabilities. It does not reason about dynamic behaviour such as timing interactions, protocol state machines, or software version-specific weaknesses, beyond what is implicitly captured through EN 50159 communication category assignments and IEC 61508 SIL classifications [7, 31]. Modelling errors or omissions in the ABox will directly affect the inferred attack surface, just as they would affect a human analyst – the quality of the assessment is bounded by the accuracy and completeness of the architectural documentation from which the ABox is generated.

These limitations do not invalidate the approach but emphasise that the ontology and rules constitute one component of a broader assurance process, to be complemented by expert review, active testing practice, and, where feasible, dynamic analysis.

5.2 Towards a Railway-Specific Methodology Adaptation

The framework demonstrated here addresses a structural gap in current railway cybersecurity practice: existing methodologies either lack railway-specific semantics or provide insufficient formal reproducibility for safety-constrained environments [1, 2, 16]. By encoding IEC 62443, EN 50159, IEC 61508, and CLC/TS 50701 as machine-interpretable SWRL rules, it transforms qualitative standards compliance into quantitative, architecture-driven scope definitions. The ETCS Level 2 demonstration produced 2 120 inferred triples from 424 ABox individuals, identifying 58 vulnerable flows, 7 EN 50159 Category 3 flows, 5 CriticalImpactTarget assets, and 5 high-risk maintenance paths – all traceable to normative sources and not derivable from any single standard alone. The framework is limited to a single ETCS Level 2 architecture and structural vulnerabilities only; validation against broader signalling families (CBTC, legacy relay-based interlocking) and real-world incident data remains as future work [2, 8]. In a companion work, the inferred outputs demonstrated here – vulnerable flows, EN 50159 threat classifications, SIL-critical attack paths, maintenance risk vectors, and Layer 4 coverage metrics – will be used to define a complete railway-specific methodology adaptation with detailed phases, activities, and artefacts. In that setting, the ontology-based reasoning will serve as the core mechanism for scope justification, threat-model generation, and quantitative coverage reporting within the pentest lifecycle – providing regulators and safety engineers with auditable, standards-grounded evidence of testing completeness [5, 9].

6 Conclusions

This paper has presented a formal OWL/SWRL-based ontological framework for automated cybersecurity scoping and prioritisation for railway signalling systems, bridging the gap between IEC 62443, EN 50159, EN 50126, and CLC/TS 50701 and the practical requirements of penetration-testing scope definition [5, 7, 9, 30].

The framework encodes domain-agnostic ICS security concepts in a three-group TBox, operationalises them through a four-layer SWRL rule architecture, and demonstrates – on a representative ETCS Level 2 architecture – that automated reasoning can derive reproducible, standards-justified penetration-testing priorities not obtainable from existing approaches. The demonstration produced 2 120 inferred triples, identifying 58 vulnerable flows, 7 EN 50159 Category 3 flows, 5 `CriticalImpactTarget` assets, and 5 high-risk maintenance paths, with all findings traceable to normative sources.

The main limitation of the present work is that it has been evaluated on only one signalling family and lacks a comprehensive ground-truth dataset of real-world vulnerabilities and test cases, so its generalisability and practical effectiveness still need

to be validated in broader industrial settings [2]. As a next step, we plan to apply the framework as a formal scoping mechanism within a railway-specific PTES adaptation (PTES-R), using the inferred vulnerable flows, EN 50159 threat classifications, SIL-critical attack paths, and coverage metrics demonstrated here to define reproducible, auditable penetration-testing phases for railway signalling systems.

References

[1] R. Kour, A. Patwardhan, A. Thaduri, R. Karim, "A review on cybersecurity in railways," *Proc. Inst. Mech. Eng. Part F*, vol. 237, no. 1, pp. 3–20, 2023. DOI: 10.1177/09544097221089389.

[2] S. Soderi, D. Masti, Y. Z. Lun, "Railway cyber-security in the era of interconnected systems: A survey," *IEEE Trans. Intell. Transp. Syst.*, vol. 24, no. 6, pp. 6764–6779, 2023. DOI: 10.1109/TITS.2023.3254442.

[3] N. Ibadah, C. Benavente-Peces, M.-O. Pahl, "Securing the future of railway systems," *Sensors*, vol. 24, no. 24, p. 8218, 2024. DOI: 10.3390/s24248218.

[4] International Electrotechnical Commission, "IEC 62443-4-1: Security for industrial automation and control systems – Part 4-1," Edition 1.0, 2018.

[5] CENELEC, "CLC/TS 50701: Railway applications – Cybersecurity," 2021.

[6] European Union, "Directive (EU) 2022/2555 (NIS2)," 2022.

[7] CENELEC, "EN 50159: Railway applications – Safety-related communication in transmission systems," 2010.

[8] J. Braband, "Towards an IT security risk assessment framework for railway automation," arXiv:1704.01175, 2017.

[9] International Electrotechnical Commission, "IEC 62443-3-2: Security risk assessment for system design," 2020.

[10] MITRE, "MITRE ATT&CK for Industrial Control Systems," `https://attack.mitre.org/matrices/ics/`, 2020.

[11] "The Penetration Testing Execution Standard (PTES)," `http://www.pentest-standard.org/`, 2014.

[12] S. Fenz, T. Neubauer, "Ontology-based information security compliance determination," *Inf. Comput. Secur.*, vol. 26, no. 5, pp. 551–567, 2018.

[13] I. Horrocks et al., "SWRL: A Semantic Web Rule Language combining OWL and RuleML," *W3C Member Submission*, 2004.

[14] B. Motik, U. Sattler, R. Studer, "Query answering for OWL-DL with rules," in *Proc. ISWC*, 2005.

[15] European Union Agency for Railways, "ERTMS/ETCS System Architecture and Interfaces," `https://www.era.europa.eu/`, 2016.

[16] A. Staves, A. Gouglidis, S. Maesschalck, D. Hutchison, "Risk-based safety scoping of adversary-centric security testing on operational technology," *Safety Sci.*, vol. 175, p. 106481, 2024.

[17] NIST, "NIST SP 800-82 Rev. 3: Guide to OT Security," 2023.

[18] ENISA, "Hands-on CLC/TS 50701: Railway applications – Cybersecurity," 2023.

[19] ISECOM, "Open Source Security Testing Methodology Manual (OSSTMM)," 2017.
[20] NIST, "NIST SP 800-115: Technical Guide to Information Security Testing," 2008.
[21] N. Maunero, F. Rosa, P. Prinetto, "Towards cybersecurity risk assessment automation: An ontological approach," in *Proc. IEEE DASC/PiCom/CBDCom*, 2023.
[22] M. J. Page et al., "The PRISMA 2020 statement," *BMJ*, vol. 372, p. n71, 2021.
[23] A. Shanley, M. N. Johnstone, "Selection of penetration testing methodologies," in *13th Australian InfoSec Mgmt. Conf.*, 2015.
[24] M. B. Imtias et al., "Comparative analysis of penetration testing frameworks," *J. Artif. Intell. Comput. Eng.*, 2025.
[25] SANS Institute, "ICS613: ICS/OT Penetration Testing & Assessments," 2026.
[26] T. Klíma, "PETA: Methodology of information systems security penetration testing," *Acta Informatica Pragensia*, 2016.
[27] M. R. Albrecht, R. B. Jensen, "The vacuity of the open source security testing methodology manual," in *NSPW*, 2020.
[28] F. de Franco Rosa, M. Jino, "A survey of security assessment ontologies," arXiv:1706.09771, 2017.
[29] International Electrotechnical Commission, "IEC 62443-3-3: System security requirements and security levels," Edition 1.0, 2013.
[30] CENELEC, "EN 50126: Railway applications – RAMS," latest edition, 2017.
[31] International Electrotechnical Commission, "IEC 61508: Functional safety of E/E/PE safety-related systems," Parts 1–7, 2010.
[32] T. L. Saaty, *The Analytic Hierarchy Process*. McGraw-Hill, 1980.
[33] B. Schneier, "Attack trees: Modeling security threats," *Dr. Dobb's J.*, vol. 24, no. 12, pp. 21–29, 1999.
[34] X. Ou, W. F. Boyer, M. A. McQueen, "MulVAL: A logic-based network security analyzer," in *Proc. 14th USENIX Security Symp.*, pp. 113–128, 2005.
[35] A. Brazhuk, "Towards automation of threat modeling based on a semantic model of attack patterns and weaknesses," *Proc. ITNT*, Dec. 2021. DOI: 10.48550/arXiv.2112.04231.
[36] A. Hosseini, W. Kastner, T. Sauter, "Ontology framework supporting security-by-design of industrial control systems," *IEEE Trans. Ind. Inform.*, 2025.
[37] J. Mozzaquatro, N. Antunes, M. Vieira, H. Madeira, "An ontology-based cybersecurity framework for the Internet of Things," *Comput. Secur.*, Sep. 2018.
[38] I. Arsuaga, M. Aguado, E. Jacob, P. Saiz, "A framework for vulnerability detection in European train control railway communications," *Security Commun. Netw.*, vol. 2018, p. 5634181, 2018. DOI: 10.1155/2018/5634181.
[39] Danish media and regulatory reports, "Ransomware incident at DSB: Supply-chain attack on Danish rail operations," 2022.
[40] J.-B. Lamy, "Owlready2: A Python library and quadstore for ontology-oriented programming," *SoftwareX*, vol. 12, p. 100539, 2020. DOI: 10.1016/j.softx.2020.100539.