

VPA: An Ontology for Verified Permissions and Authorisations

Maarten Duhoux^{1,*}, Ghislain Atemezing^{1,*}

¹European Union Agency for Railways (ERA), 120 rue Marc Lefrancq, 59300 Valenciennes, France

Abstract

Granting a permission is rarely a single administrative act: it is the outcome of a distributed process in which a requesting body files a request, accredited bodies verify evidence of compliance against legal requirements, and a permitting body finally grants the permission. In the European railway sector, this pattern underpins vehicle (type) authorisation, vehicle registration, trackside approvals and safety certification, yet the data documenting these processes has been scattered across heterogeneous registers without a shared conceptualisation. This paper presents the Verified Permissions and Authorisations (VPA) ontology, a compact OWL ontology published at <https://w3id.org/vpa> under the EUPL 1.2 license, modelling the complete lifecycle of evidence-based permission granting: requests decomposed into cases and scopes, evidence constituted by evidence documents, compliance checks against sections of legal requirements, restrictions, and time-bounded permissions. The ontology is validated through competency questions, accompanied by SHACL shapes, and is already reused in production by the ERA ontology for vehicle authorisation and registration data. We argue that VPA fills a genuine gap between policy languages such as ODRL, which express what a permission allows, and the administrative reality of how a verified permission comes into existence.

Keywords

Ontology engineering, Permissions, Vehicle authorisation, Railway data, SHACL, Linked data

1. Introduction

The European Union Agency for Railways (ERA) manages, through its legal obligations regarding registers, large datasets that document permission-based processes mandated by the Safety and Interoperability Directives.¹ In the One-Stop Shop (OSS), applicants submit so-called applications for vehicle (type) authorisations of which some key evidence is verified by Notified Bodies (NoBo), Designated Bodies (DeBo) and other assessment bodies; in the register called ERADIS, such NoBo certificates and declarations of verification are registered; in the one called ERATV, authorised vehicle types are published; and in the European Vehicle Register (EVR), vehicles are registered for operation, reusing much of the same underlying data. Despite this overlap, no common, practical ontology existed for these *verified permission* processes: each register structured its data independently, hampering interoperable exchange and forcing stakeholders — National Safety Authorities, Registering Entities, keepers, and manufacturers — to resubmit and reconcile the same facts repeatedly.

The Verified Permissions and Authorisations (VPA) ontology² was designed to close this gap. Its central insight is that a wide family of administrative processes instantiates one generalised pattern consisting of two consecutive activities: (i) a *verification process*, in which evidence regarding compliance to sections of requirements is checked by an *appropriate body*, and (ii) a *permission process*, in which a *permitting body* converts a request and its verified evidence into a permission enabling the *requesting body* to act. The pattern covers not only vehicle (type) authorisation and vehicle registration, but also trackside approval, safety certification of railway undertakings and safety authorisation of infrastructure managers, and — beyond rail — accreditations, permits and licensing schemes with a

Seventh International Workshop on Semantics for Transport and Logistics, 2026

*Corresponding author.

EMAIL: maarten.duhoux@era.europa.eu (M. Duhoux); ghislain.atemezing@era.europa.eu (G. Atemezing)

ORCID: 0009-0003-6653-6123 (M. Duhoux); 0000-0003-1562-6922 (G. Atemezing)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹Directive (EU) 2016/797, <https://eur-lex.europa.eu/eli/dir/2016/797/oj>

²<https://w3id.org/vpa>

natural bridge to W3C Verifiable Credentials [1].

The contributions of this paper are: (1) the VPA ontology (v1.2), a documented, dereferenceable OWL ontology under the EUPL 1.2 license; (2) its validation through competency questions and a companion set of SHACL shapes [2]; and (3) evidence of real-world adoption through its reuse in the ERA ontology [3] for vehicle authorisation and registration. The remainder of the paper reviews related work (Section 2), describes the ontology (Section 3), presents its validation (Section 4) and its reuse in the ERA ontology (Section 5), before concluding (Section 6).

2. State of the Art on Vehicle and Permission Ontologies

Vehicle and railway ontologies. Semantic descriptions of vehicles have been addressed from several angles. In the automotive world, vocabulary efforts around *auto.schema.org* describe vehicles as products, while VSSo [4] formalises car signals and attributes for connected-vehicle applications. In the railway domain, the RaCoOn rail core ontologies [5] pioneered cross-system data integration; the Rail Topology Ontology [6] provides a base model of rail infrastructure; and the ERA ontology [3, 7] is the reference vocabulary of the Agency’s knowledge graph, covering infrastructure, vehicle types and authorisations for route compatibility checks. These vocabularies describe the *artefacts* — vehicles, types, networks — in rich detail, but none conceptualises the administrative process by which an authority permits an actor to place such artefacts on the market or operate them.

Permission, policy and compliance ontologies. On the permission side, ODRL [8] is the W3C standard for expressing policies over assets, including permissions, prohibitions and duties; it excels at stating *what is allowed*, but — like access-control vocabularies and license-clearance frameworks for RDF resources — is agnostic to the institutional workflow that produces a legally valid permission.

The Core Criterion and Core Evidence Vocabulary (CCCEV [9]) was “designed to support the exchange of information between organisations or persons (more generally Agents) defining Requirements and organisations or persons responding to these Requirements by means of structured or unstructured Evidences”. This vocabulary aligns well with the objectives of VPA, and can be preferred if the given requirements themselves are machine processable (using exact criteria of compliance) and the underlying arguments for a compliance check can be made public. These assumptions are absent in the more abstract VPA, where a legal resource and a section therein is seen as sufficient to explain the why of a compliance check and its eventual restriction. Another difference with VPA lies in the non-expressability of many of the Requirements’ details and the underlying details of the Compliance check result: VPA does not try to document in detail which process was executed in order to arrive to a certain conclusion, as these aspects may remain confidential between the `vpa:AppropriateBody` and the `vpa:RequestingBody`. As such, VPA-classes are meant to be subClassed from if more details on this (timing, contents, etc.) are to be modeled.

The European Legislation Identifier (ELI)³ provides identifiers and an ontology for legal resources, which VPA directly reuses to anchor requirements in legislation. Finally, the W3C Verifiable Credentials data model [1] standardises cryptographically verifiable claims issued by an issuer to a holder; it shares VPA’s triangular role structure but operates at the level of credential exchange, not of compliance assessment against sectioned legal requirements.

The gap. To the best of our knowledge, no existing vocabulary captures, in one coherent model, the trio of accredited roles (requesting, appropriate, permitting body), the decomposition of a request into cases and scopes, the aggregation of evidence documents into evidence, per-section compliance checks with structured results, restrictions conditioning acceptance, and time-bounded validity of evidence and permissions. VPA occupies precisely this niche with deliberate economy, reusing ELI, FOAF, W3C ORG, SKOS and OWL-Time, and adding only the process-specific terms that were missing. At its core, it hence allows to refine further processes of approval by inheritance.

³<https://eur-lex.europa.eu/eli-register/about.html>

3. The VPA Ontology

3.1. Methodology and Publication

VPA was developed iteratively following the Linked Open Terms methodology [10]: requirements were elicited from the ERA business processes listed in Section 1, formalised as competency questions (Section 4), implemented in OWL, and evaluated at each release. The ontology is serialised in Turtle, published with content negotiation under the persistent namespace <https://w3id.org/vpa#> (prefix `vpa`), and documented with WIDOCO.⁴ Every release is traced with `skos:changeNote` entries, and superseded terms are never deleted but flagged with `owl:deprecated`, `vs:term_status` "archaic" and `dcterms:replaces` links (e.g., `vpa:documentsCase` was replaced by `vpa:documentCase` in v1.2), guaranteeing that existing data remains interpretable — a policy we consider essential for vocabularies backing legal registers.

3.2. Core Model

Figure 1 depicts the conceptual model. Version 1.2 defines 12 classes, 36 object properties (two of which survive only as deprecated terms) and 4 datatype properties in the `vpa` namespace, complemented by reused terms from the ERA, FOAF, W3C ORG, ELI, SKOS and OWL-Time vocabularies; its axiomatisation — inverse and functional properties, union domains, existential restrictions and disjointness axioms — falls within the OWL 2 DL profile. The sources, SHACL shapes and issue tracker are publicly developed in the ERA GitLab repository.⁵ Three disjoint body classes, all subclasses of `era:Body` (an organisation or person, per W3C ORG and FOAF), carry the process:

`vpa:RequestingBody` *issues* a `vpa:Request` and *requests* a `vpa:Permission`; `vpa:AppropriateBody` — an accredited assessor — *certifies* requirements and *delivers* evidence documents; and `vpa:PermittingBody`, whose organisational role *grants* the permission.

A `vpa:Request` is *constituted by* at least one `vpa:Case` (an existential restriction on `vpa:constitutedBy`), each case *concerning* one or more `vpa:Scopes` — the concrete material objects of the permission, such as a vehicle type and its area of use. Each case carries a `vpa:permissionType` drawn from a SKOS concept scheme, and both requests and cases expose a functional `vpa:status` property tracking their treatment by the permitting body.

Evidence is modelled at two levels, a distinction introduced in v0.5 after practical experience: `vpa:EvidenceDocument` (a `foaf:Document`, typically a certificate or declaration, versioned via `vpa:versionNumber` and `dcterms:replaces`) and `vpa:Evidence`, an `rdf:Bag` of evidence documents representing the collective proof submitted for a scope, case or request through the property hierarchy `vpa:supportsScope` / `vpa:supportsCase` / `vpa:submittedFor`. The presence of EvidenceDocuments in the same Evidence-Bag infers its common use and support for the Case or Request it links to: the Documents must be considered together in order to express the check and the related restrictions. No other use of `rdf:Bag` is made.

The verification activity itself is reified as `vpa:Compliance`: each check records the *checked requirement*, the *checked section* (a SKOS concept registered under the requirement), a boolean `vpa:isCompliant`, a structured `vpa:checkResult` and a free-text `vpa:statement`. Crucially, `vpa:Requirement` is a subclass of `eli:LegalResource`, so that Technical Specifications for Interoperability (TSIs) and other legal acts are referenced as first-class, dereferenceable legal resources (of which versions are available, encoded in the URI or as an available property). Partial compliance is handled by `vpa:Restriction`, attachable to a compliance check or an evidence document via `vpa:withRestriction`; a deliberate scope note forbids attaching restrictions to aggregate evidence, since only individual checks and documents are precise enough to justify one. Permissions and evidence documents are time-bounded via `vpa:valid` ranging over `time:Interval`, and carry lifecycle statuses from controlled vocabularies. Finally, `vpa:permits` links a permission to the SKOS concept

⁴<https://github.com/dgarijo/Widoco>

⁵<https://gitlab.com/era-europa-eu/public/interoperable-data-programme/era-ontology/vpa-ontology>

Table 1

Representative competency questions and the VPA terms that answer them.

Competency question	Answering terms
CQ1: Which permissions has a given body requested, and in which requests?	vpa:requests, vpa:requestFor/vpa:requestedIn
CQ2: Of which cases is a request constituted, and what is the scope and permission type of each case?	vpa:constitutedBy, vpa:concerns, vpa:permissionType
CQ3: Which evidence, and which individual evidence documents, support a given case or scope?	vpa:supportsCase, vpa:supportsScope, rdfs:member, vpa:documentCase
CQ4: Which sections of which legal requirements were checked, by whom, and with what result?	vpa:checkedRequirement, vpa:checkedSection, vpa:isCompliant, vpa:checkResult, vpa:delivers
CQ5: Under which restrictions was an evidence document accepted?	vpa:withRestriction, vpa:regarding
CQ6: During which time interval is a permission (or an evidence document) valid?	vpa:valid, time:Interval
CQ7: What is the current status of a request, its cases, and the resulting permission?	vpa:status, vpa:permissionStatus, vpa:documentStatus
CQ8: What does a granted permission actually permit its holder to do?	vpa:permits, vpa:grants

4.2. Data Validation with SHACL

OWL axioms alone cannot enforce data quality in an open-world setting, so VPA ships a companion SHACL [2] shapes graph (vpa-shapes.ttl) defining node shapes for all thirteen central classes — from era-sh:RequestShape (a request must expose its status, target permission and constituting cases) to era-sh:ComplianceShape (checked section and requirement must be IRIs of the expected classes, vpa:isCompliant must be boolean), each violation carrying a human-readable message. These shapes are executed in the ingestion pipelines of the ERA registers, turning the ontology’s implicit expectations into actionable validation reports. In addition, the ontology was iteratively checked with the OOPS! pitfall scanner [11] and standard DL reasoners for consistency, and the disjointness axioms among the three body classes and among the eight process classes guard against the most common instantiation errors.

Downstream registers refine, rather than duplicate, these generic shapes. The EVR Application Guide⁶ tightens vpa:constitutedBy, vpa:requestedIn and vpa:concerns to require an era:VehicleRegistrationCase, -Application and -Set respectively, and reuses the property path vpa:requestedIn/vpa:constitutedBy inside a SHACL-SPARQL constraint — \$this dct:type ...PRE4RP . ?cert vpa:requestedIn/vpa:constitutedBy \$this . FILTER NOT EXISTS { ?cert dct:type ...APIS } — to enforce, beyond OWL’s reach, that a case of type PRE4RP must belong to a certificate of type APIS.

This inherent refining and shaping is expected and explains why the VPA-ontology does not provide specific aspects of the verification and approval processes: validity, state, transitions between conclusions, use of the permission in processes. These aspects are to be refined in the child-ontology and in our case have been in the ERA-ontology in several processes as will be discussed now.

5. Reuse in the ERA Ontology

The strongest evidence for VPA’s adequacy is its adoption by the ERA ontology [7], the reference vocabulary of the agency’s knowledge graph. Rather than duplicating process concepts, the ERA ontology *specialises* VPA. For the vehicle (type) authorisation (VA) process:

era:VehicleAuthorisation	rdfs:subClassOf vpa:Permission .
era:VehicleTypeAuthorisation	rdfs:subClassOf vpa:Permission .

⁶<https://rinf.data.era.europa.eu/era-vocabulary/evr-appGuide/shapes.ttl>

```
era:VehicleTypeAuthorisationApplication rdfs:subClassOf vpa:Request .
era:VehicleTypeAuthorisationCase rdfs:subClassOf vpa:Case .
era:DocumentSet rdfs:subClassOf vpa:Evidence .
era:CertificationLevelDocument rdfs:subClassOf vpa:EvidenceDocument . # NoBo
era:ECDeclaration rdfs:subClassOf vpa:EvidenceDocument .
```

Here the applicant plays `vpa:RequestingBody`, the authorising entity (ERA or an NSA, per Regulation (EU) 2018/545⁷) plays `vpa:PermittingBody`, and NoBos/DeBos play `vpa:AppropriateBody`, delivering certificates whose compliance checks reference TSI sections as ELI legal resources. A comprehensive guide using the ERA and VPA ontologies for this use case is publicly available⁸.

The vehicle registration (VR) process demonstrates the payoff of the shared model. A registration application (`era:VehicleRegistrationApplication` $\sqsubseteq$ `vpa:Request`) filed by a keeper does not restate the authorisation facts: the authorisation case is linked into the registration request via `vpa:constitutes`, so tools can traverse from a registered vehicle to its authorisation, area of use, and underlying evidence — across what used to be the EVR/ERATV/ERADIS register boundaries — in a single SPARQL query. This cross-register chaining, impossible before, is the interoperability dividend the ontology was designed to deliver, and it generalises: safety certification and authorisation follow the same pattern, and the alignment of the three VPA roles with issuer–holder–verifier makes a future bridge to Verifiable Credentials [1] natural.

6. Conclusion

We presented VPA, a compact ontology for verified permissions and authorisations, published under a persistent IRI with open license, documentation, shapes and a disciplined deprecation policy. VPA is positioned between artefact-centric vehicle vocabularies and policy languages: it models how permissions *come to be*, grounded in accredited verification of evidence against sectioned legal requirements. Its validation rests on competency questions from real regulatory processes and on SHACL shapes running in production pipelines, and its reuse in the ERA ontology shows the abstraction holds where it matters — in the data of the European railway registers. Future work includes publishing the controlled vocabularies for statuses, check results and permission types as ERA reference data, extending coverage to the safety certification data streams of the NSAs, and formalising the mapping to W3C Verifiable Credentials. VPA will also be exercised with manufacturers in a machine-to-machine (M2M) test bed for efficient vehicle authorisation,⁹ supported by an open-source TypeScript library for producing and consuming ERA RDF resources.¹⁰

Acknowledgments

The authors thank the ERA register teams for their feedback on successive revisions of the ontology.

Declaration on Generative AI

The authors used generative AI assistance for editing and reviewing the manuscript text; all content was reviewed and edited by the authors, who take full responsibility for it.

⁷https://eur-lex.europa.eu/eli/reg_impl/2018/545/oj

⁸<https://eva-cld-guide-06fc8d.gitlab.io/>

⁹<https://gitlab.com/era-europa-eu/public/interoperable-data-programme/era-ontology/efficient-vehicle-authorisation/va-m2m>

¹⁰<https://gitlab.com/era-europa-eu/public/interoperable-data-programme/era-ontology/efficient-vehicle-authorisation/era-rdf-ts-resources>

References

- [1] M. Sporny, T. Thibodeau, I. Herman, M. B. Jones, G. Cohen, Verifiable credentials data model v2.0, W3C Recommendation, <https://www.w3.org/TR/vc-data-model-2.0/>, 2025.
- [2] H. Knublauch, D. Kontokostas, Shapes constraint language (SHACL), W3C Recommendation, <https://www.w3.org/TR/shacl/>, 2017.
- [3] J. A. Rojas, M. Aguado, P. Vasilopoulou, I. Velitchkov, D. V. Assche, P. Colpaert, R. Verborgh, Leveraging semantic technologies for digital interoperability in the European railway domain, in: The Semantic Web – ISWC 2021, volume 12922 of *Lecture Notes in Computer Science*, Springer, 2021, pp. 648–664. doi:10.1007/978-3-030-88361-4_38.
- [4] B. Klotz, R. Troncy, D. Wilms, C. Bonnet, VSSo: A vehicle signal and attribute ontology, in: Proceedings of the 9th International Semantic Sensor Networks Workshop (SSN 2018), co-located with ISWC 2018, volume 2213 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2018, pp. 56–63.
- [5] J. Tutchter, J. M. Easton, C. Roberts, Enabling data integration in the railway industry using RDF and OWL: The RaCoOn ontology, *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering* 3 (2017). doi:10.1061/AJRUA6.0000859.
- [6] S. Bischof, G. Schenner, Rail topology ontology: A rail infrastructure base ontology, in: The Semantic Web – ISWC 2021, volume 12922 of *Lecture Notes in Computer Science*, Springer, 2021, pp. 597–612. doi:10.1007/978-3-030-88361-4_35.
- [7] European Union Agency for Railways, The ERA ontology, <https://data-interop.era.europa.eu/era-vocabulary/>, 2024. Accessed 2026-07-17.
- [8] R. Iannella, S. Villata, ODRL information model 2.2, W3C Recommendation, <https://www.w3.org/TR/odrl-model/>, 2018.
- [9] J. D. Cock, E. Dhondt, F. Karabulut, A. Sofou, E. Stani, B. V. Nuffelen, CCCEV core criterion and core evidence vocabulary 2.2.0, SEMIC Community data model, <https://semiceu.github.io/CCCEV/releases/2.2.0/>, 2026.
- [10] M. Poveda-Villalón, A. Fernández-Izquierdo, M. Fernández-López, R. García-Castro, LOT: An industrial oriented ontology engineering framework, *Engineering Applications of Artificial Intelligence* 111 (2022) 104755. doi:10.1016/j.engappai.2022.104755.
- [11] M. Poveda-Villalón, A. Gómez-Pérez, M. C. Suárez-Figueroa, OOPS! (Ontology Pitfall Scanner!): An on-line tool for ontology evaluation, *International Journal on Semantic Web and Information Systems* 10 (2014) 7–34. doi:10.4018/ijswis.2014040102.