

# MCS v0.5.1 Companion Pack

---

*A Supervisory-Evidence Ontology for Agentic AI under EU Law: Candidate Minimum Conceptual Set and Temporal Extension*

**Author:** Jeroen Janssen, Apparens (Deventer, Netherlands)

**Date:** 25 April 2026

**Version:** v0.5.1 (release)

**License:** CC BY 4.0

**Intended use:** arXiv ancillary upload (cs.CY primary, cs.AI and cs.SE secondary). Zenodo deposit for DOI assignment.

**Contents:** the candidate Minimum Conceptual Set paper (main text), the framework specification (human-readable equivalent of the vocabulary), and the source-file appendices listed in the table of contents below.

---

## Table of contents

1. Paper (mcs\_paper\_v0\_5\_1.md): full text of the candidate Minimum Conceptual Set paper.
2. Framework specification (mcs\_framework\_specification\_v0\_5\_1.md): human-readable equivalent of the vocabulary.
3. Appendix A: README v0.5.1.
4. Appendix B: ERRATA v0.5.1.
5. Appendix C: Artefact manifest v0.5.1.
6. Appendix D: Scenarios v0.5.1.
7. Appendix E: Coding rubrics v0.2.
8. Appendix F: Validation pack v1.
9. Appendix G: Vocabulary (Turtle).
10. Appendix H: SHACL profiles v1.2 (Turtle).
11. Appendix I: SME proportionality profile v0.2 (Turtle).
12. Appendix J: Revocation model v0.2 (TLA+).
13. Appendix K: TLC config.
14. Appendix L: TLC status v0.5.

## 1. Paper

---

# A Supervisory-Evidence Ontology for Agentic AI under EU Law: Candidate Minimum Conceptual Set and Temporal Extension

---

Jeroen Janssen, April 2026

**Version:** v0.5.1 **DOI:** [10.5281/zenodo.19758441](https://doi.org/10.5281/zenodo.19758441) **License:** CC BY 4.0 **Cite as:** Janssen, J. (2026). *A Supervisory-Evidence Ontology for Agentic AI under EU Law: Candidate Minimum Conceptual Set and Temporal Extension* (v0.5.1) [Working paper]. Zenodo. <https://doi.org/10.5281/zenodo.19758441>

## Abstract

---

Agentic AI has outpaced the ontologies intended to govern it. Commercial and academic ontologies released between 2024 and 2026 cluster around a shared enterprise core of Agent, Skill, Policy, Memory, and Outcome, but none was designed to produce evidence that a European supervisor can ingest. Current supervisory practice relies on ad-hoc documentation produced per controller and per request. This paper proposes a shared representational layer for agentic AI accountability evidence under EU law, structured in three components. The first is a candidate Minimum Conceptual Set of 23 conceptual slots, separated into an agent-behaviour core (twelve slots) and a supervisory-evidence layer (eleven slots). Under strict reuse-zero accounting these 23 slots correspond to 16 net-new classes plus 7 reuse slots (5 DPV reuses in slots 6 to 10 and 2 PROV-O reuses in slots 3 and 19); the Turtle vocabulary contains 69 owl:Class declarations once subtypes, support classes, and named categories are counted. Each slot is mapped to evidence needs arising under the GDPR, the AI Act, or NIS2, or is motivated by structured reading of a twenty-five-case sample of EU ADM enforcement. The second is a temporal extension expressed in OWL-Time and made structurally checkable through SHACL shapes for delegation validity, revocation propagation records, policy versioning, and evidence decay. The third is an integration layer that reuses GDPRov, DPV, and PROV-O through owl:imports rather than reinventing their concepts. A Standardised Supervisory Ingestion Interface is identified as a research direction in Section 8 rather than claimed as delivered contribution. The paper does not claim reference-architecture status. It claims that the synthesis and design choices are designed to be defensible, reproducible, and testably better than ad-hoc practice; the empirical validation tracks remain pending. Validation is pre-registered through three open tracks; Track 1 (inter-rater consistency on the case sample) is committed to for a follow-up revision. A companion v1.2 SHACL release (`mcs_profiles_v1_2.ttl`) ships Profile A (AP-inspired permissive, quantitative) and Profile B (CNIL/German-guidance-inspired stricter, qualitative) alongside this paper, with a size-based SME proportionality profile (`mcs_sme_profile_v0_2.ttl`) as a separate axis. Limitations include empirical validation deficit shared with every framework in the reviewed corpus, a single-coder empirical base pending Track 1 execution, documented distributive effects that specification work cannot correct, and dependency on external regulatory coherence that is empirically contingent.

**Keywords:** agent ontology, EU AI Act, GDPR Article 22, supervisory evidence, SHACL, PROV-O, temporal governance, agentic AI, Minimum Conceptual Set

© 2026 J. Janssen Apparens

---

## 1. Introduction

---

A European data protection authority receives a notification that an agentic AI system at a regulated controller has produced an outcome that may engage GDPR Article 22, AI Act Annex III, or NIS2 Article 23. The supervisor opens an inquiry and asks the five standard questions. Who took this action. On what authority. Drawing on what data. Under what policy version. With what human intervention.

Today the answer arrives in one of three forms. The first is a vendor compliance dashboard, queryable only through that vendor's UI, comparable to no other controller's dashboard, and structurally invisible to the supervisor's analytical tooling. The second is a long narrative attestation, expensive, retrospective, and not machine-readable. The third is nothing structured at all: a Slack reconstruction, an email trail, a manually compiled audit memo. None of these scale. None of these are queryable. None of these support cross-controller comparison, which is what European-level supervision under the EDPB, the European AI Office, and the GDPR Article 60 cooperation procedure requires to operate at all.

The problem is not a lack of regulation. The GDPR, the AI Act, NIS2, DORA, the CRA, and the revised Product Liability Directive supply operative provisions in dense quantity. Nor is it a lack of technical ontology work: W3C PROV-O, OWL-Time, SHACL, DPV, and GDPRov have been available and stable for years. The gap sits between the two: no agent ontology released between 2024 and 2026 encodes the concepts a European supervisor needs, and no supervisory instrument specifies the ontological form in which it expects evidence to arrive. The consequence is operational, not theoretical. Supervisors spend disproportionate time on schema reconciliation before substantive analysis begins. Small and medium controllers face compliance costs they cannot absorb because every evidence request is bespoke. Data subjects face procedural opacity because the structure in which their case is represented is different at every controller.

This paper proposes a shared representational layer for agentic AI accountability evidence under EU law. The design has three layers. A core Minimum Conceptual Set (MCS) of twenty-three classes, separated into an agent-behaviour core (twelve classes) and a supervisory-evidence layer (eleven classes), each mapped to evidence needs arising under a specific provision of the GDPR, the AI Act, or

NIS2, or in structured motivation against a sampled enforcement corpus (Section 4). A temporal and validation layer expressed in OWL-Time with SHACL shapes for delegation, revocation, policy versioning, and evidence decay (Section 6). An integration layer that reuses GDPRov, DPV, and PROV-O via owl:imports and invokes sectoral regimes (MDR, MiFID II, CRR/CRD, IDD, Awb) without absorbing them (Section 4.4 and 7.6). A Standardised Supervisory Ingestion Interface (SSII), discussed as a research direction in Section 8 item 2, is not claimed as a delivered contribution at this release state; any SSII endpoint in the worked scenarios is illustrative and hypothetical.

The research questions are two. RQ1 is structural: what is the minimum conceptual set necessary to produce supervisor-ingestible accountability evidence for agentic AI under EU law. RQ2 is operational: how does the MCS map to specific statutory obligations under the GDPR, the AI Act, and NIS2, and what SHACL structures would allow supervisors to validate that evidence is structurally complete before it leaves the controller's perimeter.

The central claim is narrower than it may appear. The MCS is not the only possible synthesis. Competing syntheses are possible and will likely follow. The contribution is a candidate defensible minimum, grounded in existing open ontologies where they suffice, extended where the legal-semantic gap requires new work, targeted at supervisory-grade accountability for agentic AI operating under EU law. The test of originality is not uniqueness. It is whether the synthesis is defensible, reproducible, and testably better than ad-hoc practice for the teams that would use it. Validation is treated as the second half of the contribution, but the release-state evidence remains a validation design rather than completed validation. A companion validation pack publishes three pre-registered open tracks: inter-rater consistency on the twenty-five case sample, SHACL throughput on real agent runtimes, and structural fit across agent topologies the paper has not stress-tested.

**Positioning.** This paper is a working paper, not a reference-architecture claim. The distinction is deliberate. The MCS has a single-author empirical base pending independent inter-rater validation, an ingestion interface that remains a research direction rather than a specified contribution, and specification governance within a single commercial consultancy rather than a standards body. Each of these is an item in the research agenda of Section 8. None is resolved at release state. A reader who treats this paper as reference architecture is applying a standard the paper does not claim.

The paper has an author disclosure. The MCS was developed by the author through Apparens, a consultancy providing AI governance services. The author has a professional interest in the adoption of specifications of this class. Section 7.8 analyses the distributive consequences of that interest directly rather than deflecting the point.

---

## 2. Methods

---

### 2.1 Ontology selection

Search executed January 2024 to April 2026 across the W3C Standards directory, the ISO/IEC JTC 1/SC 42 published catalogue, the CEN-CENELEC JTC 21 work programme, the Gartner Magic Quadrant for Metadata Management Solutions (2024 and 2025), arXiv categories cs.AI, cs.MA, cs.SE, and vendor documentation pages. Search terms: "agent ontology", "AI governance ontology", "agentic AI knowledge graph", "agent specification language", "agent control protocol", "compliance ontology", "supervisory ontology", plus Dutch equivalents within the Dutch supervisory ecosystem.

Inclusion criteria, all required: active version published or updated after 1 January 2024; public schema available as downloadable OWL, RDF, SHACL, or TTL, or an academic description in a peer-reviewed venue, an arXiv preprint, or a standards-body deliverable; either at least three independent citations or development under standards-body governance.

Exclusion criteria, any one excluding: closed-source with no public schema and no academic description; fewer than three independent citations and no standards-body governance; no version after 31 December 2023; consultancy white paper without published schema or peer review.

Fourteen ontologies met the criteria and appear in Table 4.1: Aviso, Atlan, Microsoft Fabric IQ, Skan Agentic Ontology of Work, Agent Spec, ACP Framework, OntoBoom, Foundation AgenticOS, PROV-O, OWL-Time, GDPRov, DPV, ISO/IEC 42001 plus 23894 plus 5338 (grouped), and CEN-CENELEC JTC 21 drafts (prEN 18286, prEN 18228, prEN 18284).

### 2.2 Coverage matrix coding

Two matrices operate at different resolutions. The survey matrix in Section 4.1 rates 14 ontologies  $\times$  18 concepts = 252 cells. Each cell carries one of three values: present (formal class or property in a published schema; evidence is an OWL class declaration, an RDF property, a SHACL shape, or an explicit named construct in schema documentation); partial (covered conceptually in prose documentation; evidence is a recognisable named concept without a formal construct); absent (no construct or concept matching the row label returned by structured search of schema and prose). The survey matrix establishes that no single ontology suffices and that GDPR legal-semantic concepts cluster in GDPRov and DPV rather than in agent ontologies. The deep-coding matrix referenced in Appendix F.3 rates 25 MCS class slots  $\times$  5 ontologies (PROV-O, GDPRov, DPV, Agent Spec, ACP Framework) = 125 cells under a stricter three-value rubric that additionally tests owl:imports-reusability and operational-criteria preservation. The deep-coding matrix establishes which MCS classes are reusable via owl:imports versus which require new MCS-side constructs.

Coding performed on 23 April 2026 in round 1 by the author for both matrices. Round 2 self-consistency audit is scheduled for 7 May 2026 under the identical rubric published as Appendix F.2 (survey matrix)

and Appendix F.3 (deep-coding matrix). The audit is intra-rater rather than inter-rater. Intra-rater agreement does not rule out systematic coder bias shared between rounds. Track 1 of the validation pack (Appendix E.1) commissions an independent second coder; until Track 1 reports, both matrices are read as one author's structured reading rather than as inter-subjectively validated.

### 2.3 Case-law selection

Two selection criteria operate in parallel. Criterion A: CJEU judgments interpreting GDPR Article 22 directly, decided between 1 January 2023 and 31 December 2025, and cited in published DPA guidance available by April 2026. Criterion B for Appendix B: enforcement decisions (CJEU, national courts, DPAs) decided between 25 May 2018 and April 2026 in which Article 22 applicability was the contested or determined issue, or the case was cited by at least one DPA as an Article 22 reference point.

Two cases meet Criterion A: C-634/21 SCHUFA Holding (Scoring), 7 December 2023, and C-203/22 Dun and Bradstreet Austria, 27 February 2025. Twenty-five cases meet Criterion B and form the empirical validation sample in Appendix B. The two-case scope under Criterion A is acknowledged as narrow; the wider sample under Criterion B carries the empirical weight.

### 2.4 Supervisor selection

Dutch market supervisory bodies named in the AP-RDI Final advice on the supervisory structure for the AI Act (7 November 2024), plus bodies that have issued explicit AI-related guidance with publication date on or before April 2026. Nine distinct entities, ten supervisory roles: AFM, DNB, IGJ, AP, RDI, Nederlandse Arbeidsinspectie/SZW, NCSC, ILT, ACM consumer-and-markets, ACM energy. ACM is counted twice for two functionally distinct mandates. References to "ten supervisors" elsewhere should be read as "ten supervisory roles". The list should be reassessed as Article 70 designations crystallise via the Uitvoeringswet AI-verordening (expected Q4 2026).

### 2.5 MCS derivation principle

Stated principle: the MCS is the minimum set of classes such that every operative provision in GDPR Articles 4, 22, and 30, and AI Act Articles 12, 13, 14, 26, 72, and 73, has at least one ontological correlate.

Counting convention. The MCS organises 23 conceptual slots. Under strict reuse-zero accounting, 5 DPV reuses (slots 6 to 10) and 2 PROV-O reuses (slots 3 and 19) each contribute 0 net-new classes. The slot count of 23 therefore reduces to 16 net-new MCS classes: 6 in mcs-core (slots 1, 2, 4, 5, 11, 12) plus 10 in mcs-supervisory (slots 13 to 18, 20 to 23). NonQualifyingADM is the 16th net-new class, anchored in the empirical validation in Appendix B.5. The Turtle vocabulary contains 69 owl:Class declarations once subtypes, support classes, and named categories are counted; this raw figure is the file count, not the conceptual count.

A more permissive reading yields figures in two ranges. Adding sectoral subclasses (MiFID II suitability, MDR Rule 11, Awb besluit classes) extends the MCS to roughly 26 to 30 classes. Counting each subtype hierarchy as separate classes produces a range of approximately 30 to 43 classes depending on which hierarchies are expanded. The 23 slot count is the working figure for MCS scope; the 16 net-new-class figure is the working figure for ontology size; the 69 owl:Class figure is the file count. Conflating slot count with class count is the inconsistency previously noted in earlier drafts. Cardinality is justified by the operative-provision principle plus the strict reuse-zero accounting plus two documented categories of slack, not by any claim of uniqueness.

**Architectural layering.** The 23 classes are organised into two layers, separated architecturally in Section 4.3 and operationalised as distinct owl:imports targets in a planned core-plus-supervisory file split. The agent-behaviour core (mcs-core) contains twelve classes covering agent action primitives, decision and intervention semantics, and delegation-and-policy temporal structure. The supervisory-evidence layer (mcs-supervisory) contains eleven classes covering sectoral regime pointers, supervisory scope, typed log events, monitoring observations, drift indicators, evidence artefacts, serious incidents, and the NonQualifyingADM residual. The layering is primarily a coherence move: the title "agent behaviour ontology" was not honest about the scope of classes 14-23, which are supervisory scaffolding rather than agent behaviour. The layering is secondarily a deployment move: a controller deploying only the agent-behaviour core for internal auditability does not need to populate the full supervisory-evidence layer until a supervisor asks.

### 2.6 Regulation verification

Every statutory provision cited in the paper was retrieved from EUR-Lex in its authentic-language form. GDPR (Regulation 2016/679) in English and Dutch (OJ L 119, 4 May 2016). AI Act (Regulation 2024/1689) in English and Dutch (OJ L series, 12 July 2024). NIS2 Directive (2022/2555) in English (OJ L 333, 27 December 2022). Product Liability Directive (Directive 2024/2853) in English (OJ L series, 18 November 2024). Each direct quotation, paraphrase, and article reference in Section 6 and Appendix A was matched against EUR-Lex line by line.

Cross-check findings. Section 6.3 NIS2 cell: Article 23(3) limbs are alternative (disjunction), not conjunctive; severity qualifier is part of the test; the capable-of-causing branch applies alongside the caused branch; entity classification is separate from the significance threshold. AI Act Article 12(2) Dutch translation cross-checked against the Publicatieblad; no material discrepancy. GDPR Article 22(1) quotations: no discrepancies. CJEU judgments retrieved from curia.europa.eu; paragraph numbers match official texts.

### 2.7 Case-classification protocol (promoted from Appendix F.6)

The 25-case empirical sample in Appendix B is classified using a five-step protocol that is the operational heart of Finding 2 (NonQualifyingADM) and Finding 3 (Solely/Assisted boundary instability). The full rubric, including sector-calibrated thresholds and output schemas, appears in Appendix F.6. The protocol in compressed form:

*Step 1. Operative-provision test.* What GDPR or equivalent provision is the decision actually grounded on? If the operative ground is GDPR Arts 5, 6, 9, 25, or 35, or Art. 8 ECHR, rather than Art. 22: the case is classified Out-of-Art.22-Scope and populates the NonQualifyingADM structural category.

*Step 2. AI involvement test.* Was ADM or profiling present in the factual substrate? If no, the case is NonAutomated and exits the protocol.

*Step 3. WP251rev.01 meaningful-intervention test.* If ADM is present and Art. 22 is operative, does the human intervention in the record meet all six WP251rev.01 plus AP handvatten properties (hasAuthorityToDeviate, hasUnderstandingOfLogic, consideredAllRelevantData, hasMarginOfDiscretion, interventionDuration above threshold, authority actually exercised)? If yes, Assisted. If no, Solely. Any property failure is recorded individually.

*Step 4. Borderline cases.* A case is borderline when Art. 22 reasoning is present in the decision but the operative ground is Arts. 5, 6, 9, 25, or 35. Primary coding per Step 1 (Out-of-Art.22-Scope); permissive alternative per Steps 2 and 3. Both recorded.

*Step 5. Profile B cross-check.* Under Profile B (CNIL/German-guidance-inspired stricter), the threshold for “solely automated” is lower, which may reclassify borderline Assisted cases as Solely. Any reclassification recorded.

**Operationalised thresholds.** The 60-second default for interventionDuration is not statutory. It is a defensible operational presumption based on the AP handvatten requirement that intervention be “non-symbolic” combined with the working assumption that symbolic review can be performed in under one minute per decision. Sector calibration: for complex clinical triage decisions the threshold is raised to 5 minutes; for high-volume fraud-detection alerts the threshold is lowered to 15 seconds conditional on evidence that the reviewer examines case-specific features rather than rubber-stamping. The sector calibrations are declared per case, not applied silently.

**Reliability.** Round-1 coding performed by the author on 23 April 2026. Round-2 self-consistency scheduled 7 May 2026 under shuffled case order with seeded randomisation; kappa against round 1 reported as intra-rater agreement. Track 1 of the validation pack commissions an independent second coder with target Cohen’s  $\kappa \geq 0.70$  (substantial agreement per Landis and Koch 1977); failure at  $\kappa < 0.60$  triggers rubric revision rather than cell-level reconciliation.

---

### 3. Related work

#### 3.1 Surveyed ontologies (operational context)

Vendor ontologies model enterprise workflow fluently but ignore data-protection primitives. Atlan leads the 2025 Gartner Magic Quadrant for Metadata Management and offers a context graph for agentic AI with decision-trace reification, but does not import DPV and has no public reference schema. Aviso’s Ontology Layer has paying enterprise customers but no public schema. Microsoft Fabric IQ has one named pilot (ENMAX Power) and is in public preview. OntoBoom offers SHACL/OWL authoring but not regulatory semantics. None of these ships a public reference schema at the time of writing.

Academic and framework-level ontologies have different gaps. Agent Spec (Open Agent Specification) represents a declarative, framework-agnostic approach emphasising portability through symbolic references. The portability creates a tension with supervisory requirements: an agent too portable fragments its audit trail across runtimes, which impedes longitudinal reconstruction. The ACP Framework (Agent Control Protocol) offers the strongest delegation semantics in the surveyed set, with verifiable chained delegation, cryptographic signatures, mandatory expiry on delegated authority, and transitive revocation propagating across delegation chains. ACP does not encode GDPR legal primitives.

Open normative scaffolds remain PROV-O (W3C Recommendation, 30 April 2013) and OWL-Time (W3C Recommendation, 2017; Candidate Recommendation update, 2022). Two GDPR-specific extensions bridge PROV-O to data-protection semantics: GDPRov (Pandit and Lewis, SEMANTiCS 2017) extends PROV-O and P-Plan for consent and data-lifecycle provenance; GDPRtEXT provides a vocabulary for GDPR text structure allowing direct linkage of ontology concepts to specific articles and recitals. The Data Privacy Vocabulary (DPV), published by the W3C Data Privacy Vocabularies and Controls Community Group as a Final Community Group Report (v2.3, 25 February 2026), is the most mature open vocabulary for data-protection concepts. Neither DPV nor GDPRov carries formal regulatory authority. DPV is a Community Group Report, not a W3C Recommendation. Neither has been endorsed by the EDPB, the European Commission, or any national DPA as an authoritative representation of GDPR semantics. Integration of DPV and GDPRov into the MCS rests on technical maturity and community adoption, not on regulatory status.

ISO/IEC JTC 1/SC 42 has produced prose standards (22989 terminology, 23053 ML framework, 23894 risk management, 5338 lifecycle, 42001 management system, 42005 impact assessment, 42006 certification bodies, 38507 governance) but no machine-readable ontology. CEN-CENELEC JTC 21 is drafting harmonised standards under Standardisation Request M/593 plus M/613 to provide presumption of conformity with the AI Act. prEN 18286 (AI quality management system) reached public Enquiry stage 30 October 2025 to 22 January 2026. prEN 18228 (risk management) and prEN 18284 (dataset governance) are identifiable by scope but technical text is confidential.

### 3.2 Computational law and legal informatics (intellectual context)

The intellectual ancestors of the MCS are eight works across forty years of computational law and AI-and-law scholarship. Sergot et al. (1986), The British Nationality Act as a logic program, is the canonical demonstration that a statute can be formalised and executed. The MCS borrows the insight that operative statutory provisions can be rendered as machine-readable predicates; it does not attempt to formalise entire statutes as logic programs. Sartor (2005) supplies the foundational framework for defeasible and deontic reasoning; the dual SHACL profiles in Section 5.5 (AP-inspired permissive, CNIL/German-guidance-inspired stricter) are a concrete application of defeasibility to Article 22 classification. Bench-Capon and Sartor (2003) on case-based reasoning with theories and values is the methodological ancestor of the empirical-validation procedure in Appendix B. Pagallo (2013) on legal personhood of agents is the reason the MCS does not attempt to give agents separate legal standing; it encodes agent action as attributable to human or organisational principals via the Legal Role Triad. Verheij (2003, 2017) on argumentative disagreement explains the structural form that supervisory divergence takes in the MCS: two SHACL profiles, same underlying evidence graph. Hashmi et al. (2018) on compliance representation in BPM predates but parallels the approach; the MCS differs by targeting agentic rather than process-model compliance. Atkinson et al. (2020) supplies the explanation-type taxonomy that the HumanIntervention class makes auditable. Pandit et al. (2019) is the foundational DPV report on which five MCS classes depend by direct import.

Position summary. The MCS is not novel in believing that statutes admit ontological representation, that legal rules are defeasible, that case-law is empirically tractable as a validation set, that legal personhood of agents is contested, that supervisory divergence is structurally similar to argumentative disagreement, that compliance representation has a pre-history in BPM, that explanations have a typology, or that DPV is the working data-protection vocabulary. What the MCS adds, against this background, is a specific operational target: a 23-class, SHACL-validatable, supervisor-ingestible evidentiary fabric for agentic AI under EU law, anchored in named statutory provisions and validated against named enforcement cases. This is a narrower and more applied contribution than any of the eight works above, and it presupposes them.

### 3.3 Position against closest contemporary competitors

Three contemporary efforts sit closer to the MCS in problem selection and require direct positioning.

*Cobbe, Singh and Sheridan (2025), "Governance and accountability frameworks for AI agents" (working paper)*. Closest competitor on supervisor-ingestibility. Treats the same problem at a higher level of abstraction, arguing for accountability-by-design principles without committing to specific ontological structures or regulatory anchors. Differentiation: the MCS commits to named GDPR, AI Act, and NIS2 provisions and to SHACL-validatable evidence; Cobbe et al. stop at principles. The MCS presupposes their principles and operationalises them against EU statutory text.

*Agent Spec (Open Agent Specification)*. Closest competitor on cross-runtime portability. Achieves portability through symbolic references that decouple agent specifications from specific runtimes. Differentiation: the MCS sacrifices portability for evidentiary grounding. An agent running under Agent Spec is portable across runtimes; the cost is that its audit trail fragments, which defeats the longitudinal supervisory reconstruction the MCS targets. The two projects solve adjacent but orthogonal problems; direct comparison on a single axis is category-confused.

*ACP Framework (Agent Control Protocol)*. Closest competitor on delegation semantics. Provides verifiable chained delegation with cryptographic signatures, mandatory expiry, and transitive revocation. Differentiation: ACP's delegation model is stronger than the MCS's DelegationGrant class in cryptographic guarantee, but ACP does not encode GDPR legal primitives (data subject, purpose, lawful basis, consent lifecycle) and cannot represent the three-way Decision taxonomy or the Article 22 drawsStronglyOn predicate. An implementation that uses ACP for delegation and the MCS for legal-semantic ingestion is architecturally available; the MCS provides the owl:imports surface and neither project objects to the composition.

No competing synthesis the author has reviewed achieves the combination of GDPR legal-semantic anchoring, three-way Decision taxonomy with drawsStronglyOn predicate, dual-profile SHACL for supervisory divergence, and AI Act typed log events. The combination is the claim, not any component in isolation. If a reader knows of a synthesis that achieves the combination with fewer classes, closer statutory fit, or better empirical anchoring against the same 25-case sample, the author commits to engaging it in a revised edition.

---

## 4. RQ1 answer: a Minimum Conceptual Set for supervisory evidence

---

### 4.1 Cross-ontology coverage

Table 4.1 maps fourteen surveyed ontologies against eighteen concepts a European supervisor plausibly needs. For readability in portrait layout, the matrix is presented as four sub-tables grouped by ontology family: 4.1a vendor ontologies; 4.1b framework ontologies; 4.1c open normative ontologies; 4.1d standards-body drafts. The 252-cell total (14 × 18) is preserved across the four sub-tables. Legend: ● present (formal class or property in a published schema); P partial (covered conceptually or in standards-body prose); ○ absent. ISO/IEC and CEN-CENELEC JTC 21 drafts are prose standards without machine-readable schemas; under the strict legend, their cells are P or ○ rather than ●, regardless of how explicitly a concept is named in the prose.

Column header abbreviations: Av = Aviso; At = Atlan; FIQ = Microsoft Fabric IQ; Skan = Skan Agentic Ontology of Work; AS = Agent Spec; ACP = Agent Control Protocol Framework; OB = OntoBoom; FAOS

= Foundation AgenticOS; PROV = PROV-O; OWL-T = OWL-Time; GDPRov = GDPRov; DPV = Data Privacy Vocabulary; ISO = ISO/IEC 42001 + 23894 + 5338 (grouped); JTC21 = CEN-CENELEC JTC 21 drafts (prEN 18286, 18228, 18284).

**Table 4.1a. Cross-ontology coverage, vendor ontologies (4 × 18 cells).**

Concept	Av	At	FIQ	Skat
Agent	P	P	P	●
Action / tool invocation	P	P	P	●
Memory read	P	●	P	●
Memory write	P	●	P	●
Delegation	○	○	○	P
Revocation / invalidation	○	P	○	○
Policy / constraint	P	●	●	●
Human oversight	P	P	P	●
Temporal interval	○	P	P	○
Authority scope	○	●	P	P
Provenance	○	●	P	P
Versioning	○	P	P	P
Risk classification	○	○	○	P
Data subject reference	○	○	○	○
Purpose of processing	○	P	○	P
Lawful basis	○	○	○	○
Personal-data cat. (Art. 9)	○	○	○	○
Consent record (lifecycle)	○	○	○	○

**Table 4.1b. Cross-ontology coverage, framework ontologies (4 × 18 cells).**

Concept	AS	ACP	OB	FAOS
Agent	●	●	P	●
Action / tool invocation	●	P	P	●
Memory read	P	○	○	P
Memory write	P	○	○	P
Delegation	P	●	○	P
Revocation / invalidation	○	●	○	○
Policy / constraint	●	P	●	●
Human oversight	P	○	○	P
Temporal interval	○	●	P	○
Authority scope	P	●	○	P
Provenance	●	P	○	P
Versioning	●	○	P	○
Risk classification	○	○	○	P
Data subject reference	○	○	○	○
Purpose of processing	P	○	○	P
Lawful basis	○	○	○	○
Personal-data cat. (Art. 9)	○	○	○	○
Consent record (lifecycle)	○	○	○	○

**Table 4.1c. Cross-ontology coverage, open normative ontologies (4 × 18 cells).**

Concept	PROV	OWL-T	GDPRov	DPV
Agent	●	○	●	●
Action / tool invocation	●	○	●	P
Memory read	P	○	○	○
Memory write	P	○	○	○
Delegation	●	○	P	P
Revocation / invalidation	●	○	●	●
Policy / constraint	○	○	P	●
Human oversight	○	○	○	P
Temporal interval	P	●	P	P
Authority scope	P	○	○	P
Provenance	●	○	●	P
Versioning	P	○	P	P
Risk classification	○	○	○	P
Data subject reference	○	○	●	●

Purpose of processing	○	○	●	●
Lawful basis	○	○	●	●
Personal-data cat. (Art. 9)	○	○	P	●
Consent record (lifecycle)	○	○	●	●

**Table 4.1d. Cross-ontology coverage, standards-body drafts (2 × 18 cells).**

Concept	ISO	JTC21
Agent	P	P
Action / tool invocation	P	P
Memory read	○	○
Memory write	○	○
Delegation	P	P
Revocation / invalidation	P	P
Policy / constraint	P	P
Human oversight	P	P
Temporal interval	P	P
Authority scope	P	P
Provenance	P	P
Versioning	P	P
Risk classification	P	P
Data subject reference	○	P
Purpose of processing	P	P
Lawful basis	○	P
Personal-data cat. (Art. 9)	○	P
Consent record (lifecycle)	○	P

Cell ratings reflect round-1 coding on 23 April 2026 by a single coder; round-2 self-consistency audit is scheduled for 7 May 2026 under the rubric in Appendix F.2. The matrix is one coder's structured reading, not inter-subjectively validated. Three cells carry parenthetical refinements: Revocation in GDPRov and DPV is scoped to consent revocation rather than general provenance invalidation; Risk classification in DPV is via the DPV-RISK extension module rather than DPV core; Temporal interval in DPV is via the DPV-TIME extension module (dpv:Duration and related constructs) rather than DPV core.

Five findings follow from the matrix. ACP Framework provides the strongest delegation semantics with formal validity windows, transitive revocation, and cryptographic verification, exceeding PROV-O's basic delegation support. Agent Spec advances portability at the cost of supervisory depth: its abstraction layer enables cross-platform deployment but fragments audit trails. PROV-O remains the normative spine with cross-domain, open, formally specified coverage of Agent, Activity, Delegation, Association, Revision and Invalidation, but its absence of Policy, Purpose, Data Subject, Risk, and Authority Scope makes it necessary but not sufficient. Four of the five GDPR-supervisory-critical concepts (Data Subject, Purpose of Processing, Lawful Basis, Personal-Data Categories, Consent Record) are present in GDPRov and DPV but absent in every agent ontology surveyed; integration, not invention, closes this gap. Risk classification, intended purpose, sectoral regime, supervisory scope, typed log events, monitoring observations, drift indicators, supervisory evidence artifacts, and serious incident are absent everywhere except partially in ISO/IEC and JTC 21 drafts in prose, which are not machine-readable.

Three gaps remain genuine after integration with the existing open ontologies. First, no agent ontology encodes AI Act-specific legal categories (Provider, Deployer, Manufacturer, Authorised Representative, Importer, Distributor) with first-class properties. Second, no ontology carries first-class Delegation Validity Window, Revocation Propagation Semantics, Evidence Refresh, or Policy Version Successor relations. Third, supervisory divergence on Article 22 is not representable as configurable profiles over the same underlying evidence graph.

## 4.2 Integration with GDPRov, DPV, and the Legal Role Triad

The integration pattern reuses existing open vocabularies via owl:imports rather than reinventing their concepts. A minimal bridging pattern connects an agent action to its GDPR legal context:

```

gov:AgentAction a owl:Class ;
  rdfs:subClassOf prov:Activity , gdprov:DataProcessingStep .

gov:AgentActionShape a sh:NodeShape ;
  sh:targetClass gov:AgentAction ;
  sh:property [
    sh:path dpv:hasLegalBasis ;
    sh:minCount 1 ;
    sh:class dpv:LegalBasis ;
    sh:message "Art. 6(1) GDPR requires a lawful basis for every processing activity"
  ] ;
  sh:property [
    sh:path dpv:hasPurpose ;
    sh:minCount 1 ;
    sh:class dpv:Purpose ;
    sh:message "Art. 5(1)(b) GDPR: purpose limitation"
  ] ;
  sh:property [

```

```

sh:path dpv:hasDataSubject ;
sh:class dpv:DataSubject ;
sh:message "Data subject must be identifiable for Art. 15 access requests"
] .

```

This pattern reduces the absent gap from nine MCS concepts to zero for the GDPR legal-semantic layer, without inventing new classes. The remaining gap is the AI Act legal-role layer (Provider, Deployer, Manufacturer) and the three-way Decision taxonomy, both of which require new classes because DPV predates the AI Act.

The Legal Role Triad ontology declares Provider, Deployer, Distributor, Importer, Manufacturer, and AuthorisedRepresentative as subclasses of a new gov:AIActRole root. A single organisation can simultaneously hold dpv:Controller under GDPR and gov:Deployer under the AI Act; the gov:RoleAssignment class (subclass of prov:Attribution) makes role attribution time-indexed so that a supervisor can reconstruct role ownership at any past moment without conflating the two statutory frames.

### 4.3 The 23 MCS slots, split into core and supervisory layers

The MCS is organised into 23 conceptual slots collectively yielding 16 net-new classes under the strict reuse-zero accounting in Section 2.5, separated architecturally into mcs-core (12 slots, 6 net-new classes) and mcs-supervisory (11 slots, 10 net-new classes). The separation keeps the title scope-honest: classes covering sectoral regime pointers, supervisory scope, and evidence artefacts are supervisory scaffolding rather than agent behaviour, and are grouped separately. Class definitions appear in the companion Turtle vocabulary `mcs_vocabulary.ttl`, with a planned split into `mcs_core.ttl` and `mcs_supervisory.ttl`; the specification document `mcs_framework_specification.md` is the human-readable equivalent.

**Table 4.3a. mcs-core: agent-behaviour layer (12 classes).**

Slot	Class (MCS)	Operational definition	Regulatory anchor	Coverage in PROV-O / GDPRov / DPV / Agent Spec / ACP
1	gov:Agent with subtypes SoftwareAgent, HumanAgent, OrganizationalAgent	Supervisory-scope agent. Abstract over human, software, and organisational agents.	GDPR Art. 4(7)-(8); AI Act Art. 3(3)-(4); SCHUFA para 48	PROV-O present; others partial
2	gov:LegalRole with eight AI Act and GDPR subtypes	Role-typed attribution; time-indexed via gov:RoleAssignment	GDPR Arts 4(7)-(8), 26, 28; AI Act Art. 3(3)-(8)	DPV present for GDPR roles; AI Act roles absent everywhere
3	gov:Activity	Agent action; inherits startedAtTime, endedAtTime, wasAssociatedWith	AI Act Art. 12	PROV-O present
4	gov:ToolInvocation	Subclass of Activity; properties hasTool, hasInvocationInput, hasInvocationOutput	AI Act Art. 13	Agent Spec partial
5	gov:MemoryEvent with subclasses MemoryRead, MemoryWrite	Provenance of model/agent memory access	AI Act Art. 12; Art. 72	Atlan and Skan present; others partial
6	dpv:DataSubject (reused)	Data subject reference	GDPR Art. 4(1), 15, 22	DPV and GDPRov present; agent ontologies absent
7	dpv:PersonalData (reused)	Personal-data category hierarchy including Art. 9 special categories	GDPR Art. 9, 10, 30(1)(c)	DPV present; others absent
8	dpv:Purpose (reused)	Purpose of processing	GDPR Art. 5(1)(b), 6(4), 30	DPV present
9	dpv:LegalBasis (reused)	Lawful basis for processing	GDPR Art. 6, 9	DPV present
10	dpv:Consent plus gdprov consent lifecycle (reused)	Consent record with lifecycle	GDPR Art. 6(1)(a), 7	DPV and GDPRov present
11	gov:Decision with three subtypes (SolelyAutomatedDecision, AssistedDecisionWithMeaningfulHumanIntervention, NonAutomatedDecision), plus gov:Score and gov:drawsStronglyOn predicate	Three-way Article 22 taxonomy plus upstream attribution	GDPR Art. 22(1); SCHUFA para 73; Dun and	Absent everywhere

			Bradstreet para 40	
12	gov:HumanIntervention	Six-property meaningful-intervention schema: hasAuthorityToDeviate, hasUnderstandingOfLogic, consideredAllRelevantData, hasMarginOfDiscretion, interventionDurationSeconds, authorityActuallyExercised (per Uber/Ola Hof Amsterdam)	GDPR Art. 22(3); AI Act Art. 14; WP251rev.01; AP handvatten July 2025	Absent everywhere

**Table 4.3b. mcs-supervisory: supervisory-evidence layer (11 classes).**

Slot	Class (MCS)	Operational definition	Regulatory anchor	Coverage in PROV-O / GDPRov / DPV / Agent Spec / ACP
13	gov:Policy, gov:PolicyVersion	Versioned policy with temporal validity and supersession	AI Act Art. 9 risk management; Art. 17 QMS	Atlan, Skan, Agent Spec, FAOS present
14	gov:RiskClassification with subclasses ProhibitedAI, HighRiskAI, LimitedRiskAI, MinimalRiskAI	Risk classification aligned with AI Act	AI Act Art. 5, 6, 50, Annex III; ISO/IEC 23894	ISO/IEC partial (prose); absent as formal class elsewhere
15	gov:IntendedPurpose	Developer-attested property of the AI system distinct from dpv:Purpose	AI Act Art. 3(12), 13	JTC 21 partial (prose); Agent Spec Intent partial
16	gov:SectoralRegime with instances MiFID II, MDR, CRR_CRD, IDD, AWB	Sectoral regime pointers with competentSupervisor links	Sectoral statutes	Absent everywhere
17	gov:SupervisoryScope with instances for ten Dutch supervisors plus EU bodies	Competent authority with defined supervisory mandate	AI Act Art. 70; AP-RDI 7 November 2024	Absent everywhere
18	gov:DelegationGrant with subclass of prov:Delegation plus gov:RevocationEvent, gov:ValidityInterval, gov:AuthorityScope	Temporal class with validity windows and revocation semantics	GDPR Art. 28; AI Act Art. 25	ACP Framework present; PROV-O partial (basic)
19	gov:ProvenanceChain (integration point)	Reuse of prov:wasDerivedFrom, wasInformedBy, wasInfluencedBy	AI Act Art. 10; Art. 12	PROV-O present
20	gov:LogEvent with subclasses RiskIdentificationLogEvent, PostMarketMonitoringLogEvent, OperationMonitoringLogEvent, BiometricIdentificationLogEvent, IncidentLogEvent	Typed log events aligned with Article 12(2)-(3) and NIS2 Art. 23	AI Act Art. 12(2)-(3); Art. 26(6); NIS2 Art. 23	Generic logging ubiquitous; typed-to-stature subclasses absent everywhere
21	gov:MonitoringObservation	Subclass of prov:Entity with observedMetric, observedAtTime, monitoringPurpose	AI Act Art. 72; ECB Guide to Internal Models; SAFEST	Absent everywhere
22	gov:DriftIndicator with subclasses DataDrift, ConceptDrift, ModelPerformanceDrift, ControlDrift	Drift properties driftMagnitude, driftDetectedAt, triggersRefresh	AI Act Art. 72 continuous performance analysis; PCAOB AS 1105	Absent everywhere
23	gov:SupervisoryEvidenceArtifact plus gov:Incident (serious) plus gov:NonQualifyingADM with five subclasses (Art5FairnessADM, Art6LawfulBasisADM, Art9SpecialCategoryADM, Art25ByDesignADM, Art35DPIAADM)	Evidence artefact with source hierarchy; serious incident class; residual class for ADM enforcement decided on GDPR provisions other than Art. 22	AI Act Art. 26(10); AI Act Art. 73; NIS2 Art. 23; GDPR Arts 5(1)(a), 6, 9, 25, 35; Art. 8 ECHR	Absent everywhere

Slot 23 counts three thematic classes (SupervisoryEvidenceArtifact, Incident, NonQualifyingADM) as one slot under the counting convention because NonQualifyingADM is the empirically-derived residual that makes the supervisory layer complete; the other two are SupervisoryEvidenceArtifact infrastructure (sources, weights, refresh) and Incident (AI Act Art. 73, NIS2 Art. 23 reporting). Under the alternative convention that counts each thematic class separately, slot 23 resolves to three classes and the total becomes 25.

The 23 figure counts under the Section 2.5 convention. Under alternative counting conventions the figure ranges from approximately 30 to 43 classes.

#### 4.4 Integrating GDPRov, DPV and the Legal Role Triad

Five of the 23 MCS classes (DataSubject, PersonalDataCategory, PurposeOfProcessing, LawfulBasis, ConsentRecord) are imported directly from DPV via owl:imports. This is documented in the Turtle

vocabulary and is the basis for the Section 7.6 observation that the GDPR legal-semantic layer is an integration problem rather than a research problem. The MCS extends DPV in two directions. First, by layering AI Act-specific concepts via the W3C EU-AIAct extension (Community Group Final Report, v2.3, 25 February 2026) and by adding the SCHUFA-grounded three-way Decision taxonomy. Second, by adding runtime-monitorable classes (MonitoringObservation, DriftIndicator, typed LogEvent) that fall outside DPV's scope of data-protection concepts proper. The relationship is integration plus extension, not replacement.

An independently arrived-at synthesis, undertaken under the alternative labels Supervisory-Grade Agent Ontology and Minimum Supervisory Ontology, decomposes the same problem into core, temporal-governance, and supervisory-evidence layers that map closely onto the MCS, the gov: extension, and the SSII split adopted here. The author treats the convergence as design corroboration rather than endorsement; the synthesis source's citation base is non-verifiable and the SGAO/MSO naming is not adopted.

#### 4.5 The EU-AIAct extension

The W3C DPVCG EU-AIAct extension (v2.3) covers Provider, Deployer, RiskLevel, CEMarking, AISystemPerformance, AIRegulatorySandbox, ProviderHumanOversightMeasure, DeployerHumanOversightMeasure, QualityManagementSystem, and PostMarketMonitoringSystem. Suggested prefix `eu-aiact`. Adopting the extension closes roughly ten of the twenty-three MCS gaps simultaneously at the AI Act legal-role layer. The extension does not cover the NonQualifyingADM residual class (GDPR Arts 5, 6, 9, 25, 35 governance paths) and does not cover the typed Article 12 log events at statutory-subclass granularity. These remain MCS-side work.

#### 4.6 Empirical validation findings (promoted from Appendix B.5)

The 25-case sample described in Section 2.3 and listed in Appendix B yields four findings that bear directly on the design of the MCS. They are promoted into the main body here because they ground specific class-level design decisions: Finding 2 grounds NonQualifyingADM (slot 23); Finding 3 grounds the dual-profile SHACL architecture in Section 5.5; Finding 4 grounds the drawsStronglyOn predicate and its dual operationalisation in v1.2.

**Finding 1: Solely dominates the litigated sample.** Fifteen of 25 cases classify as Solely under the primary coding (including borderline-Solely, post-appeal Solely, and Hamburg DPA interpretive-Solely), three as Assisted, six as Out-of-Art.22-Scope (plus one Pending). The Solely dominance is not representative of the deployed ADM population; it is a selection-bias artefact. Cases that reach courts or DPAs are cases where controllers or data subjects contested classification; selection pressure favours Solely findings because it is the classification that engages Article 22 substantively. The empirical design of the MCS must not infer population-level ADM distribution from litigated-sample distribution.

**Finding 2: ADM enforcement decided on GDPR provisions other than Article 22 is a structural category.** Five of 25 cases (B17 IB grading, B18 Clearview multi-jurisdictional, B19 live FR in schools, B20 Belastingdienst Toeslagen, B24 CNIL Clearview fine) plus B22 (SyRI, decided on Art. 8 ECHR) are Out-of-Art.22-Scope: the ADM substrate was present but the case was decided on Art. 5 fairness, Art. 6 lawful basis, Art. 9 special-category data, Art. 25 by-design, or Art. 8 ECHR proportionality. The B18 Clearview cluster and B24 CNIL Clearview fine concern the same factual complex at different DPA levels and are coded as separate enforcement instances. This is the empirical anchor for the gov:NonQualifyingADM MCS class (Section 4.3). The 24 percent share (six of 25) is the central estimate under primary coding. The case sample carries six cases with the strict all-caps BORDERLINE flag (B7, B9, B15, B16, B17, B22) plus one case (B14) with a Profile-B-conditional borderline flag, seven cases under the inclusive reading (see `mcs_case_sample_v03.csv`). Sensitivity analysis is constructed from these flags as follows. Under the F2-relevant subset flip (B7, B9, B15, B17, B22): low (6 minus 2)/25 = 16% to high (6 plus 3)/25 = 36%, yielding [16%, 36%]. Under the all-seven inclusive-BORDERLINE flip (B7, B9, B14, B15, B16, B17, B22): low (6 minus 2)/25 = 16% to high (6 plus 5)/25 = 44%, yielding [16%, 44%]. The Wilson 95% CI is [11.5%, 43.4%] anchored on k=6, n=25. The specific proportion should not be cited as a stable empirical fact; the finding that the category exists and is structurally significant is robust.

**Finding 3: The Solely/Assisted boundary is jurisprudentially unstable.** The Uber/Ola Amsterdam case sequence (B9, B10, B11, B12, B16) is the empirical confirmation that the same fact pattern can be classified differently depending on whether courts apply a formal or substantive test for meaningful human intervention. The Amsterdam District Court at first instance (B16) accepted "whole team" deciding meant Art. 22 did not apply; the Amsterdam Court of Appeal (B10, ECLI:NL:GHAMS:2023:793) reversed, holding Krakow-office reviewers' authority and competence were not established and the review was "not much more than a purely symbolic act". This finding validates the Section 5.5 design choice to operationalise AP-inspired permissive vs CNIL/German-guidance-inspired stricter divergence as two distinct SHACL profiles rather than a single classification rule. It also grounds the inclusion of gov:authorityActuallyExercised as a sixth HumanIntervention property (Table 4.3a slot 12), which operationalises the "symbolic act" distinction rather than collapsing it into hasAuthorityToDeviate.

**Finding 4: SCHUFA upstream attribution is operative in national courts.** Cases B1, B2, and B3 confirm the drawsStronglyOn path is being applied not only by the CJEU but by national courts post-SCHUFA. The Wiesbaden Verwaltungsgericht's January 2026 ruling on the SCHUFA TeamBank scorecard treated the existence of a specialised scoring model as itself evidence of strong reliance, without requiring quantification. This validates the Decision.drawsStronglyOn predicate (Section 4.3) but reveals DPAs and courts apply it qualitatively, not quantitatively. The v1.2 release (`mcs_profiles_v1_2.ttl`) operationalises this finding by making the drawsStronglyOn method a profile-level choice: Profile A (AP-inspired permissive) uses the quantitative flip-rate operationalisation as controller-side compliance instrument; Profile B (CNIL/German-guidance-inspired stricter) uses three

qualitative judicial indicia (existence of specialised scoring model; refusal frequency  $\geq 0.90$  per SCHUFA “in almost all cases”; absence of documented intervention per D&B “lacked any manual oversight”). The dual operationalisation resolves what would otherwise remain a qualitative/quantitative hedge.

Reliability caveats. All four findings rest on single-coder classification by the author in round 1 (23 April 2026). Round-2 self-consistency audit is scheduled 7 May 2026. Track 1 of the validation pack (Appendix E.1) commissions an independent second coder. Until Track 1 reports Cohen’s  $\kappa \geq 0.60$ , the findings should be treated as one coder’s structured reading of the enforcement corpus rather than as inter-subjectively validated empirical claims.

---

## 5. RQ2 answer: operational mapping to European legal categories

---

### 5.1 Article 22 doctrinal position

Two CJEU judgments frame Article 22 interpretation as of April 2026. In C-634/21 SCHUFA Holding (Scoring), judgment of 7 December 2023, the Grand Chamber held that an automated probability score is itself a “decision” within the scope of Article 22(1) where the downstream actor draws strongly on it, even if a nominal human is the formal decision-maker. Paragraph 73 establishes the upstream-attribution logic; paragraph 48 grounds the “in almost all cases” reliance test in the factual record of the SCHUFA case. In C-203/22 Dun and Bradstreet Austria, judgment of 27 February 2025, the Court clarified that “meaningful information about the logic” under Article 15(1)(h) must be functional rather than merely algorithmic: the data subject is entitled to an explanation they can act on, not a description of weight matrices.

Supervisory practice diverges. The Dutch Autoriteit Persoonsgegevens (Advies artikel 22 AVG en geautomatiseerde selectie-instrumenten, 10 October 2024, and Handvatten betekenisvolle menselijke tussenkomst, July 2025) reads the Article 22(2)(a) contractual-necessity exemption and the Article 22(3) safeguards as permitting automated selection followed by meaningful human treatment to fall outside Article 22 engagement. CNIL, AEPD (March 2024 guidance), and several German state DPAs read SCHUFA more broadly: a score drawn strongly on by a downstream decision engages Article 22(1) regardless of intervening human review unless the intervention meets a higher bar on authority and discretion. The UK ICO (post-Brexit, cited for cross-channel practice) aligns with the stricter reading: review must be “active, not a token gesture.”

The MCS does not pre-decide the divergence. Section 5.5 operationalises it as two SHACL profiles over the same underlying evidence graph. A reader who prefers the permissive reading runs Profile A; a reader who prefers the strict reading runs Profile B. An agent operating in both jurisdictions is validated against both profiles and surfaces the delta as a governance artefact.

A separate doctrinal debate concerns the legal character of Article 22 itself. One view (WP251rev.01, reinforced by SCHUFA) treats Article 22 as a general prohibition triggering controller duties regardless of data-subject invocation. A competing view (Wachter, Mittelstadt, and Floridi 2017; Veale and Edwards 2018) argues Article 22 is under-enforced as an individual right and should be reframed as a continuous procedural-accountability obligation on controllers. The current MCS design splits the difference by requiring both the three-way Decision classification (accountability side) and explicit HumanIntervention records (rights side). This hedging doubles the logging burden; the choice is legitimately open and should be made explicitly by each controller based on risk appetite and supervisory exposure.

### 5.2 Ten-supervisor model (Dutch scope)

The Dutch AI governance structure operates through ten supervisory roles held by nine distinct legal entities. Each role carries distinct, dated authorities.

Financial supervision: AFM (Autoriteit Financiële Markten), AFM Agenda 2026 (19 January 2026); DNB (De Nederlandsche Bank), General principles for responsible AI in the financial sector (25 July 2019, SAFEST framework).

Health and education: IGJ (Inspectie Gezondheidszorg en Jeugd), MDR Rule 11 guidance of 10 February 2025 classifying autonomous clinical-decision agents as Class IIa/IIb Software as a Medical Device.

Data protection and algorithm coordination: AP (Autoriteit Persoonsgegevens), coordinating algorithm supervisor since February 2025 and market supervisor for prohibited AI practices and high-risk applications; Advies artikel 22 AVG (10 October 2024) and Handvatten (July 2025). RDI (Autoriteit Digitale Infrastructuur), co-signatory on AP-RDI Final advice on the supervisory structure for the AI Act (7 November 2024); AI regulatory sandbox supervisor (mandatory as of August 2026).

Employment: Nederlandse Arbeidsinspectie (SZW), Algorithmic management: insight and workplace implications (20 November 2024).

Cybersecurity: NCSC (Nationaal Cyber Security Centrum), NIS2 implementation for AI systems in critical infrastructure.

Transport: ILT (Inspectie Leefomgeving en Transport), autonomous vehicles and transport AI.

Consumer and markets: ACM (Autoriteit Consument en Markt), competition aspects of AI platforms and algorithmic ranking systems.

Energy: ACM (energy division), smart grid AI and energy allocation algorithms.

This ten-role model creates jurisdictional complexity where a single agentic AI system can fall under multiple supervisors simultaneously. The Uitvoeringswet AI-verordening (expected Q4 2026) will clarify coordination mechanisms; current practice operates on the AP-RDI November 2024 advice.

### 5.3 Mapping matrix with inline formal sketches

Rows are agent-behaviour primitives. Columns are legal categories. Each cell gives a conditional classification rule with source and confidence grade (H/M/L), prefixed with a status flag indicating epistemic status.

Status flags. **[SD] statutory derivation**: the proposition follows directly from statutory text, recital, or settled CJEU jurisprudence. **[IP] interpretive proposal**: the proposition is a defensible reading of the source texts but rests on an interpretive step on which competent authorities or courts could differ.

**[RC] research conjecture**: the proposition extends source materials beyond what they address; no statutory provision, recital, or jurisprudence directly grounds it. Cells flagged RC are research positions, not derivable legal mappings.

The full 8 by 6 matrix is in Appendix A.3 (Table A.1). Representative cells with inline formal sketches:

#### F1 Decision × Art. 22 (statutory derivation; SCHUFA paras 50, 73).

```
Vd. Decision(d) →
  Art22Engaged(d) ↔
  ( ( SolelyAutomated(d)
    V ∃s. Score(s) ∧ drawsStronglyOn(d, s) )
    ∧ HasLegalOrSignificantEffect(d, subject) )

SolelyAutomated(d) ↔
¬∃hi. HumanIntervention(hi) ∧ hi.contributesTo(d) ∧ MeaningfulIntervention(hi)

MeaningfulIntervention(hi) ↔
hi.hasAuthorityToDeviate = true
∧ hi.hasUnderstandingOfLogic = true
∧ hi.hasMarginOfDiscretion = true
∧ hi.consideredAllRelevantData = true
∧ hi.authorityActuallyExercised = true [Profile B]
```

Captures Art. 22(1) scope plus SCHUFA upstream-attribution extension; the Art. 22(1) legal-or-significant-effect qualifier; WP251rev.01 meaningful-intervention criteria as conjunctive; the authorityActuallyExercised property responds to Uber/Ola (B10). Does not capture Art. 22(2) exemptions, Art. 22(4) special-category restriction, or Art. 22(3) safeguards when engaged under Art. 22(2)(a) or (c).

#### F2 Score × SCHUFA upstream (statutory derivation; SCHUFA paras 48, 73; Wiesbaden VG January 2026).

```
Vs, d. Score(s) ∧ Decision(d) ∧ d.used(s) →
drawsStronglyOn(d, s) ↔
PROFILE-DEPENDENT:
Profile A (AP-inspired permissive): FlipRate(s, d, ε=0.10) > M=0.50
Profile B (CNIL/German-guidance-inspired stricter): QualitativeIndicium(s, d), where
QualitativeIndicium ↔
usesSpecialisedScoringModel(s)
V refusalFrequencyGivenScore(s, d) ≥ 0.90
V hasDocumentedIntervention(d) = false
```

Captures the SCHUFA upstream-attribution path as profile-dependent; the AP-CNIL profile divergence as different operationalisations rather than different parameter pairings within a single method; the quantitative interpretation as controller-side compliance instrument; the qualitative interpretation as judicial-style application consistent with Wiesbaden VG. Does not capture cases where the score is used outside the SCHUFA factual context of “insufficient score leads to contractual refusal in almost all cases”.

#### F4 Tool invocation × NIS2 significant incident (interpretive proposal; NIS2 Article 23(3); Recital 101). Heaviest-load sketch; careless reading could have direct compliance consequences.

```
Vi. Incident(i) →
SignificantIncident(i) ↔
( ( CausedSevereOperationalDisruption(i, entity)
  V CapableOfCausingSevereOperationalDisruption(i, entity)
  V CausedSevereFinancialLoss(i, entity)
  V CapableOfCausingSevereFinancialLoss(i, entity) )
  V
  ( CausedConsiderableDamageToOthers(i)
  V CapableOfCausingConsiderableDamageToOthers(i) )
)
where Severe(i) is judged by:
ExtentOfFunctionImpaired(i) ∧
DurationOf(i) ∧
NumberOfRecipientsAffected(i) (Recital 101)

EntityClassification(entity) ∈ {Essential, Important}
determines reporting regime under Art. 3 and the
reporting cadence under Art. 23(4),
not the significance threshold itself.
```

A naive encoding as F(operational\_disruption ∧ essential\_services) under-classifies reportable incidents in three ways: conjunction would require an incident to harm both the entity and external persons

before being reportable, whereas Art. 23(3) requires either; dropping the severity qualifier lowers the threshold below statutory text; treating “essential services” as a threshold component rather than as an entity classifier conflates two distinct provisions. A controller relying on such an encoding could fail to report incidents that NIS2 Art. 23(3) requires to be reported.

The remaining seven formal sketches (F3 tool-chain propagation, F5 memory-contribution conjecture, F6 delegation Art. 28 processor chain, F7 policy lex mitior conjecture, F8 evidence AS 1105 reliability hierarchy, F9 human-intervention authority utilisation, F10 risk Annex III classification), plus three NIS2-parity sketches (F11 delegation revocation × NIS2 incident window, F12 memory write × NIS2 Art. 23(4) reporting cadence, F13 policy version change × NIS2 Art. 21 measures), appear in Appendix A.4 through A.19.

The matrix is designed to make the epistemic status of each cell explicit so that a deployer, supervisor, or reader can see where the MCS derives legal mappings and where it extends them. A controller relying on [RC]-flagged cells does so on their own risk profile. A supervisor invited to issue a binding decision on an [RC]-flagged cell should treat that as a research proposal rather than as settled law.

#### 5.4 Formal-methods scope

The MCS adopts PROV-O, OWL-Time and SHACL as the constrained baseline for general supervisory ingestion. Fuller temporal-logic verification (LTL, CTL model checking, Event-B refinement,  $\mu$ -calculus over delegation transition systems) may be appropriate for safety-critical AI subsystems where the cost of a wrong classification justifies the verification overhead. The evidence base does not support a general claim that such verification is operationally realistic for broad supervisory ingestion. Section 8 research-agenda item 6 (causal blast radius algorithms for revocation propagation) is the natural site where formal verification is most likely to earn its cost; outside that niche the MCS treats SHACL-time validation as the operative tool.

#### 5.5 Supervisory divergence as conflicting SHACL profiles (v1.2: both profiles shipped)

The mapping matrix treats Article 22 as a single target. In practice DPAs apply divergent readings of SCHUFA and Dun and Bradstreet, which an agent operating across jurisdictions encounters as conflicting compliance constraints. The MCS accommodates this by expressing divergence as two distinct SHACL profiles over the same underlying gov:Decision graph. In the companion SHACL release (mcs\_profiles\_v1\_2.ttl, released concurrently with this paper), both Profile A and Profile B ship, plus Profile C (UK ICO post-Brexit) and Profile D (Italian Garante post-Law 132/2025) as stubs pending jurisdictional case-law review.

**Profile A (AP-aligned, permissive, quantitative).** A preparatory gov:Score is treated as in-scope under Art. 22(1) only where a downstream gov:Decision with hasLegalEffect true satisfies drawsStronglyOn under the quantitative flip-rate test ( $\epsilon = 10\%$  of population range,  $M = 50\%$  flip rate) and no gov:AssistedDecisionWithMeaningfulHumanIntervention pattern intervenes. Automated selection followed by meaningful human treatment falls outside Art. 22 (AP advies, October 2024; AP handvatten, July 2025).

**Profile B (CNIL-aligned, strict, qualitative).** A preparatory gov:Score is treated as in-scope whenever a downstream decision drawsStronglyOn it under any of three judicial indicia (specialised scoring model; refusal frequency  $\geq 0.90$ ; absence of documented intervention authority), unless intervention meets all five conjunctive properties including gov:authorityActuallyExercised. CNIL guidance, AEPD March 2024, Wiesbaden VG January 2026, and several German state DPAs align with this reading.

Turtle sketches (full shapes in mcs\_profiles\_v1\_2.ttl):

```
gov:ProfileA a gov:SupervisoryProfile ;
  gov:hasDrawsStronglyOnMethod gov:QuantitativeFlipRate ;
  gov:quantitativeEpsilon "0.10"^^xsd:decimal ;
  gov:quantitativeM "0.50"^^xsd:decimal .

gov:ProfileA_Art22ScopeShape a sh:NodeShape ;
  sh:targetClass gov:Score ;
  sh:sparql [ sh:select """
    SELECT $this WHERE {
      $this a gov:Score .
      ?d gov:drawsStronglyOn $this ;
        gov:hasLegalEffect true .
      FILTER NOT EXISTS {
        ?d a gov:AssistedDecisionWithMeaningfulHumanIntervention .
      }
    } """ ] .

gov:ProfileB a gov:SupervisoryProfile ;
  gov:hasDrawsStronglyOnMethod gov:QualitativeJudicialIndicia .

gov:ProfileB_Art22ScopeShape a sh:NodeShape ;
  sh:targetClass gov:Score ;
  sh:sparql [ sh:select """
    SELECT $this WHERE {
      $this a gov:Score .
      ?d gov:drawsStronglyOn $this ;
        gov:hasLegalEffect true .
      FILTER NOT EXISTS {
        ?d gov:hasHumanIntervention ?hi .
        ?hi gov:hasAuthorityToDeviate true ;
          gov:hasMarginOfDiscretion true ;
          gov:hasUnderstandingOfLogic true ;
          gov:consideredAllRelevantData true ;
          gov:authorityActuallyExercised true .
        }
      } """ ] .
```

```
} "" ] .
```

An agent validated against both profiles surfaces the delta as a governance artefact. Decisions that pass Profile A but fail Profile B trigger a `gov:ProfileDelta` record per `gov:CrossProfileGovernanceShape`; the delta requires either jurisdictional routing, elevation of human intervention attributes, or a documented supervisory dialogue. This converts a latent regulatory conflict into an explicit decision for the controller, recorded as provenance. It does not resolve the underlying legal divergence; it makes the divergence operable.

Empirical support for the dual-profile design comes from the Uber/Ola Amsterdam case sequence (Appendix B cases B9, B10, B11, B12, B16) as analysed in Finding 3 (Section 4.6). The same fact pattern flipped classification depending on whether the court applied a formal or substantive test for meaningful intervention. The dual-profile architecture makes this jurisprudential reality representable rather than suppressing it. The inclusion of `gov:authorityActuallyExercised` as a sixth `HumanIntervention` property, operationalising the Krakow “symbolic act” distinction, is the specific design move that responds to the Hof Amsterdam finding.

## 5.6 SME proportionality profile

The dual-profile architecture of §5.5 addresses jurisdictional divergence. It does not address size-based capability asymmetry. As documented in §7.8, SMEs subject to the AI Act high-risk regime face compliance costs that larger controllers absorb as marginal extension of existing infrastructure. A specification that raises the floor of what counts as adequate evidence amplifies that asymmetry unless it carries a structural SME response.

The MCS specifies an SME proportionality profile, released as `mcs_sme_profile_v0_2.ttl` alongside the dual-jurisdictional profiles. The mechanism operates on two axes.

*Axis 1: eligibility.* An SME profile applies to a controller that meets the European Commission SME definition in EU Recommendation 2003/361/EC: fewer than 250 employees, and either annual turnover  $\leq$  €50 million or annual balance sheet total  $\leq$  €43 million. Eligibility is re-assessed annually; a controller that crosses the threshold in consecutive years exits the profile. The specific threshold is a policy choice, not an MCS-internal derivation; the €50m / €43m figures are adopted because they match the existing SME definition that national and EU instruments already use. A Member State applying a stricter threshold (for example, turnover  $\leq$  €10m for “micro” SMEs) supplies its own `gov:smeEligibilityProfile` instance pointing at national transposition text.

*Axis 2: substitution structure.* For an eligible controller, a subset of the SHACL shapes is marked `gov:SMEExempt true`. Under SME proportionality, a structured accountability narrative substitutes for the SHACL-checked structural evidence that those shapes would otherwise require. The narrative is itself structured: it must declare for each exempted shape (i) what evidence would have populated the shape if non-exempt; (ii) why the controller’s operational scale makes SHACL-checked structural evidence production disproportionate; (iii) what alternative accountability mechanism the controller provides in its place (for example, a documented manual review protocol).

The prose-substitution set is deliberately narrow. What stays in SHACL-validated scope for SMEs: the core Decision classification (`gov:DecisionShape`), the meaningful-intervention schema (`gov:HumanInterventionShape`), policy version recording (`gov:PolicyVersionShape`), delegation and revocation (`gov:DelegationGrantShape`, `gov:RevocationEventShape`), and the Article 22 scope shapes (both profiles). What becomes prose-substitutable: the Article 12 typed `LogEvent` subclasses at full statutory-subclass granularity (SMEs may log at a coarser granularity with a declared mapping); the post-market monitoring machinery (`gov:MonitoringObservationShape`, `gov:DriftIndicatorShape`) at sub-monthly cadence; the evidence decay full Bayesian structure (SMEs may declare scalar priors instead); the cross-profile governance shape (SMEs running a single jurisdictional profile may substitute a structured narrative for the delta-surfacing obligation for decisions within that single jurisdiction).

The trade-off is explicit. A controller running under SME proportionality cannot be supervised at machine-readable granularity on the prose-substituted dimensions. A supervisor who requires machine-readable evidence on a prose-substituted dimension (for example, during an incident investigation under NIS2 Art. 23(3)) may suspend the proportionality treatment for the specific dimension for the specific investigation window. The suspension is itself recorded as a `gov:ExemptionSuspension` event with a start time, an end time, a scope specification, and a supervisor identifier. Outside the suspension window the proportionality arrangement resumes automatically.

What the SME proportionality profile does not do. It does not reduce substantive legal obligation: an SME remains subject to the GDPR Art. 22, the AI Act Arts 12-14 and 26, and NIS2 Art. 21 and 23 at the same substantive standard as a larger controller. It reduces the evidentiary machinery that the MCS specifies as the default, not the underlying legal duty. It is therefore compatible with the AI Act Art. 62-63 SME support measures (regulatory sandboxes, fee reductions) rather than a substitute for them.

Turtle sketch, with full shapes in `mcs_sme_profile_v0_2.ttl`:

```
# Legacy URI retained for backward compatibility; label/comment use proportionality terminology.
gov:SMEExemptionProfile a gov:SupervisoryProfile ;
  rdfs:label "SME proportionality profile"@en ;
  gov:hasEligibilityCriterion gov:EC_SME_Definition_2003_361 ;
  gov:hasSubstitutionMechanism gov:StructuredProseNarrative .

gov:EC_SME_Definition_2003_361 a gov:EligibilityCriterion ;
  gov:maxEmployeeCount "250"^^xsd:integer ;
  gov:maxAnnualTurnover "50000000"^^xsd:decimal ;
  gov:maxBalanceSheetTotal "43000000"^^xsd:decimal ;
  gov:eligibilityReassessmentCadence "1Y"^^xsd:duration .
```

```

gov:ExemptionSuspension a owl:Class ;
  rdfs:subClassOf prov:Activity ;
  rdfs:comment "Supervisor-initiated suspension of SME proportionality for a specified dimension,
scope, and window." .

```

The framework is specified at structural level; the policy choices (specific thresholds, per-sector calibration, the set of shapes eligible for prose substitution) are controller-configurable and regulator-reviewable. The author has a commercial interest in the adoption of compliance-evidence frameworks including this one, declared in §1 and analysed in §7.8. That interest applies with particular force to SME proportionality mechanisms, because a structural SME response reduces the political economy friction against adoption by SMEs, which enlarges the market for compliance intermediation. A regulator designing the actual exemption eligibility and scope should treat the structural framework as a defensible starting point, not as a neutral recommendation.

## 6. Temporal governance

### 6.1 Revocation propagation

When a delegation is revoked, the MCS specifies five effects that a revocation event may produce, any or all of which apply: `BlocksNewWork` (no new tasks accepted under the revoked delegation), `CancelsQueuedWork` (queued tasks not yet started are cancelled), `StopsInflightWork` (in-progress tasks are stopped), `RequiresReview` (tasks completed recently must be reviewed before acceptance), `MarksArtifactsStale` (artefacts produced under the revoked authority are marked stale for downstream consumers).

The default applied in the absence of explicit configuration is `BlocksNewWork` plus `MarksArtifactsStale`. `StopsInflightWork` is an operationally risky default because it can leave consequential actions half-executed; the MCS specifies it as opt-in rather than default. A controller that adopts a different default records the divergence and justifies it; the MCS treats undocumented divergence as itself supervisory-relevant.

The property `gov:executedUnderRevokedAuthority` (boolean, default false) flags decisions produced under authority that was revoked before execution completed. Setting it to true requires explicit policy authorisation. The property `gov:cachedUnderRevokedAuthority` flags artefacts produced before revocation but accessed after, which remain lawfully produced but may need review depending on the consumer.

### 6.2 Evidence decay

The temporal reliability of evidence decays. A DPIA signed two years ago, a model validation run five years ago, and a fairness audit last quarter do not carry equivalent probative weight today. The MCS encodes decay via the property `gov:probativeWeight` (decimal in [0.0, 1.0]) and the evidence-source hierarchy `SystemLogSource`, `HumanAttestationSource`, `ExternalReportSource`, `ThirdPartyAuditSource`. Reliability weights are attached to source types following the PCAOB AS 1105 / ISA 500 / ISAE 3000 hierarchy.

**Bayesian decay with declared priors.** The MCS replaces the scalar `gov:probativeWeight` default with a Bayesian decay function anchored in declared sectoral priors. The decay takes the form:

$$\text{gov:probativeWeight}(e, t) = \text{gov:initialWeight}(e) \times \exp(-\lambda(\text{sector}) \times \Delta t)$$

where  $\Delta t$  is the elapsed time since  $e$ 's last refresh and  $\lambda(\text{sector})$  is a sector-specific decay rate expressed as a half-life. Declared priors:

Sector / regime	Half-life	Anchor
ISAE 3402 Type 2 audit evidence	12 months	ISAE 3402 standard annual re-assertion cycle
ECB model validation (CRR/CRD)	12 months	ECB Guide to Internal Models annual validation requirement
MDR post-market surveillance	6 months	MDR Annex III Section 1.1(a) periodic safety update reporting
AI Act Art. 12 automatic log retention	6 months	AI Act Art. 26(6) minimum retention floor, used as decay floor not ceiling
AI Act Art. 72 post-market monitoring	12 months	AI Act Art. 72 continuous performance analysis annual cycle
NIS2 Art. 21 measure attestation	6 months	NIS2 Art. 20(1) management-body review cadence interpreted as attestation half-life
Manual human attestation	3 months	Operational presumption pending empirical anchor

The declared priors are defensible starting points, not statutory requirements. A controller running with different priors records the divergence and its justification; the SHACL shape `gov:EvidenceDecayShape` validates that priors are either the declared defaults or a documented divergence. The priors are empirically contingent and should be re-assessed against drift-event rate data from deployed high-risk AI systems as it becomes available (Section 8 research-agenda item 8).

The AI Act Article 26(6) six-month log retention requirement provides a lower bound on storage horizon. In the Bayesian decay framework it is interpreted as a decay floor: `probativeWeight` computed

from the half-life function is clamped to  $\max(\text{Bayesian\_decay}(e, t), 0.5)$  during the first six months after creation, preventing log evidence from being treated as already-decayed within its statutory retention window. After the six-month floor, the Bayesian function applies without clamping.

The declared-priors framework does not resolve the underlying empirical anchor gap. It converts the gap from an unreasoned absence to a reasoned presumption that can be falsified as empirical data arrives. This is the standard pattern for Bayesian priors in the absence of frequentist data: declared, defensible, updatable.

### 6.3 Temporal extension with SHACL

The temporal extension defines six classes and twenty-one SHACL shapes. Classes: gov:DelegationGrant (subclass of prov:Delegation), gov:RevocationEvent (subclass of prov:Activity), gov:ValidityInterval (subclass of time:ProperInterval), gov:AuthorityScope, gov:PolicyVersion (subclass of prov:Entity), gov:SupervisoryEvidenceArtifact (subclass of prov:Entity).

Shapes (fourteen core shapes plus seven profile-layer shapes in v1.2, mcs\_profiles\_v1\_2.ttl): DelegationGrantShape (validity times, authority scope), PolicyVersionShape (version info, no overlap with predecessor), EvidenceDecayShape (refresh timestamp, 6-month justification trigger), EvidenceQualityShape (AS 1105 attributes present), MonitoringObservationShape, DriftIndicatorShape, AgentActionShape (lawful basis, purpose, data subject), RoleAssignmentShape (time-bounded), BiometricLogShape (Art. 12(3) mandatory fields), DecisionShape (three-way taxonomy well-formed), SchufaUpstreamShape (drawsStronglyOn predicates populated when the shape target is gov:Score), HumanInterventionShape (six conjunctive properties of slot 12 declared), RevocationEventShape (prov:wasInformedBy plus prov:startedAtTime), and LogEventSubclassGranularityShape (typed at AI Act statutory-subclass granularity).

The shapes fire as structural completeness checks, not as legal-compliance checks. A SHACL validation report with sh:conforms=true indicates only that the data meets the structural constraints defined in the shapes. It does not establish compliance with the AI Act, the GDPR, NIS2, DORA, the CRA, or any other instrument of Union or national law. Structural completeness is necessary but not sufficient for supervisory documentation obligations. This disclaimer is embedded in the vocabulary and shapes headers of the release artefacts.

### 6.4 drawsStronglyOn operationalisation

Both SHACL profiles in Section 5.5 reference the drawsStronglyOn predicate. The MCS operationalises it per profile: Profile A uses a quantitative instantiation parameterised by ( $\epsilon$ , M); Profile B uses a qualitative instantiation aligned with national-court practice.

*Profile A quantitative.* drawsStronglyOn(d, score,  $\epsilon$ , M) holds iff the empirical flip rate of (score, d) under counterfactual perturbation magnitude  $\epsilon$  exceeds threshold M. Computable from PROV-O DAGs via the procedure: for each score s used by decision d, perturb s by  $\epsilon$  across its empirical distribution; re-run the decision pipeline; measure the proportion of perturbation replicas in which d flips outcome. Parameters:  $\epsilon$  = 10% of population range, M = 50% flip rate. Controller-side compliance instrument; may be rejected by supervisors applying qualitative indicia.

*Profile B qualitative.* drawsStronglyOn(d, s) holds iff any of three judicial indicia applies. Indicum 1: the score uses a specialised scoring model (Wiesbaden VG January 2026 treated this as itself evidence of strong reliance). Indicum 2: the refusal frequency given insufficient score in the controller's historical decision population exceeds 0.90 (operationalises SCHUFA paragraph 48 "in almost all cases" factual finding; threshold pending jurisprudential refinement). Indicum 3: the decision lacks any documented human intervention with the five Profile B conjunctive properties (operationalises D&B paragraph 40 "lacked any manual oversight").

The dual operationalisation treats the qualitative and quantitative methods as profile-level design choices, with Profile A committing to the quantitative method as the AP-inspired permissive reading's most defensible operationalisation, and Profile B committing to the qualitative method as the CNIL/German-guidance-inspired stricter reading's most defensible operationalisation. A controller running both profiles populates both the quantitative properties (gov:flipRateAtEpsilon10) and the qualitative properties (gov:usesSpecialisedScoringModel, gov:refusalFrequencyGivenScore, gov:hasDocumentedIntervention) on each gov:Decision. SHACL fires the appropriate test per profile.

The national-court application of drawsStronglyOn is operative post-SCHUFA. The Wiesbaden Verwaltungsgericht (January 2026) treated the existence of a specialised scoring model as itself evidence of strong reliance without requiring quantification. This confirms the predicate is being applied by courts but reveals they apply it qualitatively, not quantitatively. Profile B as the qualitative-aligned profile is not a research conjecture but an alignment with decided practice.

Empirical validation of Profile A's quantitative pairings remains Section 8 research-agenda item 4: do ( $\epsilon$  = 10%, M = 50%) applied at controller-side predict the Solely versus Assisted classifications that supervisors and courts subsequently endorse. Validation requires controller participation in disclosing flip-rate measurements alongside ex-post supervisory determinations, cross-controller aggregation, and statistical analysis against endorsed outcomes.

---

## 7. Discussion

### 7.0 Cross-cutting design tensions

Six tensions shape the MCS and are resolved implicitly by specific design choices. Surfacing them lets a

reader who disagrees with the position locate the design lever rather than re-deriving it from scratch.

*Expressiveness vs auditability.* Rich semantic models capture nuance but become harder to audit consistently. Position: constrained expressiveness. The MCS records what materially affects supervision and excludes operational concepts that do not. Twenty-three classes split into twelve-class core plus eleven-class supervisory layer is the working compromise.

*Ontology richness vs supervisor usability.* Many concepts improve operational intelligence but overwhelm investigators. Position: layered design. The mcs-core is the minimum supervisory core; mcs-supervisory adds the scaffolding; richer optional layers (sectoral extensions, controller-internal concepts) sit outside the core.

*Temporal precision vs implementation feasibility.* Fine-grained validity windows and version pinning improve point-in-time reconstruction but increase storage, query, and audit cost. Position: minimum mandatory time fields with optional richer temporal logic for safety-critical niches.

*Legal caution vs operational utility.* Legal uncertainty discourages strong classification rules; operations want deterministic labels. Position: conditional rules with status flags (SD / IP / RC) and confidence grades. The MCS never asserts legal classification; it produces evidence that controllers and supervisors then classify.

*Provenance completeness vs privacy minimisation.* Full PROV-O traces strengthen evidence but multiply personal-data retention in tension with GDPR Art. 5(1)(c). Position: the MCS records what is required; it does not specify the pseudonymisation pattern that reconciles the tension. This is an open research item (Section 8 item 16).

*Single supervisor profile vs cross-jurisdictional portability.* A single SHACL profile ships faster; multiple profiles require coordination. Position: v1.2 ships Profile A (AP-inspired permissive, quantitative) and Profile B (CNIL/German-guidance-inspired stricter, qualitative) alongside this paper; Profiles C (UK ICO) and D (Italian Garante post-Law 132/2025) are stubs pending jurisdictional case-law review.

### **7.1 Internal coherence as necessary but not sufficient**

The MCS has been stress-tested internally via seven worked-example scenarios (in the companion scenarios document). All seven admit MCS class population from realistic event streams. Three (retail banking, hospital triage, municipal benefits) produce single-profile classifications under both Profile A and Profile B. Platform delivery produces a profile divergence that the dual-profile architecture was designed to surface. Revocation propagation surfaces operational control. Cross-border one-stop-shop exercises the dual-profile architecture inside GDPR Article 60 cooperation. NIS2 plus GDPR concurrent incident exercises the Appendix A.7 F4 sketch together with the F11 sketch.

Internal-consistency validation is necessary but not sufficient. The scenarios confirm the MCS admits population from realistic event streams; they do not establish that deployed controllers will populate the structure correctly, that supervisors will ingest the resulting evidence, or that the evidence will change regulatory outcomes. Those are the validation-pack questions.

### **7.2 No empirical validation on real systems**

The MCS has not been populated by any live controller, ingested by any live supervisor, or tested against any live enforcement outcome. The validation pack companion document is the programmatic response: three pre-registered open tracks on inter-rater consistency, SHACL throughput on real runtimes, and structural fit across topologies the working paper has not stress-tested. The pack is published as a companion; this version of the paper cannot claim validation from tracks not yet executed.

Track 1 execution is committed to for a follow-up revision of the paper. The validation pack is open from publication date and accepts volunteer submissions on a rolling basis; the  $\kappa$  result will be published in the next paper revision after the first qualifying submission arrives. Failure at  $\kappa < 0.60$  triggers rubric revision rather than cell-level reconciliation; publication of the failure is part of the commitment.

This limitation is shared with every prescriptive governance framework in the reviewed corpus. ISO/IEC 42001, NIST AI RMF, AgentSpec, MI9, and the academic ontology proposals surveyed in Section 3 face the same empirical deficit. The gap will not be resolved by any single instrument; it requires a research programme.

### **7.3 Structural completeness is not legal compliance**

A SHACL validation report with `sh:conforms=true` indicates only structural completeness. It does not establish compliance with the AI Act, the GDPR, NIS2, DORA, the CRA, or any other instrument. This disclaimer is embedded in the release vocabulary and shapes files. A supervisor who reads only the SHACL conforms flag without reading the underlying evidence is not substituting MCS validation for legal judgement; they are misusing the tool. The disclaimer is not self-enforcing; supervisors and deployers who do not read it are not bound by it.

### **7.4 Dual-profile implementation (v1.2 status)**

v1.2 SHACL shapes (`mcs_profiles_v1_2.ttl`) implement Profile A (AP-inspired permissive, quantitative) and Profile B (CNIL/German-guidance-inspired stricter, qualitative) over the shared `gov:SupervisoryProfile` abstract class. Profile C (UK ICO post-Brexit) and Profile D (Italian Garante post-Law 132/2025) ship as stubs. A controller running v1.2 shapes against a decision graph populates both profiles and surfaces any delta via `gov:CrossProfileGovernanceShape` and `gov:ProfileDeltaShape`. The cross-jurisdictional design is operational with v1.2 shapes.

What the dual-profile implementation does not do: resolve the underlying legal divergence; provide an authoritative choice between AP-inspired permissive and CNIL/German-guidance-inspired stricter readings; substitute for supervisor judgement on individual decisions. The implementation makes the divergence representable and auditable, which is the MCS's scope.

## 7.5 Multi-stakeholder supply chains

The MCS addresses third-party risk through the DelegationGrant class, RoleAssignment time-indexing, and the SectoralRegime class. It treats the model provider, orchestration platform, tool provider, and end-user deployer as distinct legal roles with distinct RoleAssignments against a shared Agent. The architecture reflects a shared responsibility model rather than a pure vendor-management model. It does not specify the governance interfaces between entities that would allow, for example, a model provider's QMS evidence (AI Act Art. 17) to flow structurally into a deployer's conformity assessment (Art. 43). That interface specification is a natural extension of S3 and is flagged in Section 8 as an open item.

## 7.6 Interoperability and the limits of the regulatory-coherence assumption

The MCS v1 treats the European regulatory landscape as coherent enough to permit a single structured evidence ontology to serve as a bridge across multiple regimes. A framework inventory (24 April 2026) makes the empirical support for this assumption concrete but also reveals three limits.

*NIS2 transposition is incomplete.* As of late 2025, approximately 14 of 27 EU Member States have fully transposed NIS2. The European Commission issued reasoned opinions to 19 Member States on 7 May 2025. Some Member States (Italy, Belgium, Ireland, France, Czech Republic) have transposed with variations in scope; others have not yet enacted primary transposition legislation. The Commission proposed targeted NIS2 amendments on 20 January 2026, including simplifications, in parallel with ongoing infringement proceedings. The implication for the MCS is that gov:SectoralRegime, which treats NIS2 as a uniform cross-Member-State regime, is not fully realised in practice. MCS-structured evidence can serve as input to NIS2 compliance; the claim that such evidence is automatically acceptable across Member States is weaker than the v1 default implies.

*Supervisor guidance increasingly says "existing frameworks apply".* The pattern is documented across at least four major supervisors: FCA (UK) "AI and the FCA: our approach" (2024 to 2026) is technology-neutral, no new AI-specific rules, existing SMCR, Consumer Duty, and operational resilience rules apply; BaFin (DE) guidance of 18 December 2025 positions AI as a DORA ICT asset rather than a separate regulatory category; CNIL (FR) AI recommendations (2023 to 2025) operationalise GDPR without creating new substantive obligations; AP (NL) algorithmic supervision role is grounded in existing GDPR and sectoral powers. The pattern is convergent. The MCS must therefore be positioned as complementary to existing frameworks rather than as a substitute or standalone assurance mechanism.

*Italy is a counter-example.* Law No. 132/2025 (25 September 2025, entered into force 10 October 2025) adds sector-specific AI obligations beyond the AI Act (healthcare, labour, justice, public administration). If other Member States follow Italy's template, the landscape becomes more heterogeneous at the national level while remaining coherent at the EU level. The MCS responds with Profile D as a stub in `mcs_profiles_v1_2.ttl` and with a commitment to populate the first national-variant instance in `mcs_framework_bridges.ttl` (Italy first, UK second) in a follow-up revision.

Four consequences for the MCS design follow. The MCS must document, as a limitation, that its structural coherence depends on external regulatory coherence which is empirically contingent. The bridge artefact must use explicit negative assertions (`bridge:hasNoMappingTo`) where a framework is absent, withdrawn, or incomplete. The MCS should publish a 12-month refresh cadence with interim-trigger-based delta notes. The MCS should resist the temptation to add AI-specific concepts where existing sectoral or horizontal concepts (GDPR, DORA, MDR, UK GDPR) suffice.

## 7.7 Dual-use considerations

The released artefacts contain only data-model components: an RDF/OWL vocabulary covering 23 conceptual slots and contributing 16 net-new classes under strict reuse-zero accounting (69 owl:Class declarations in the Turtle file when subtypes, support classes, and named instances are counted), twenty-one SHACL shapes (fourteen core temporal-extension shapes including three SME-supporting shapes added in v0.5.1, plus seven profile-layer shapes in v1.2), three further shapes in the SME proportionality profile, a workload generator, and a framework-inventory reconnaissance document. These artefacts on their own cannot observe, gate, or enforce any behaviour. The concern arises from the design trajectory: the full MCS contemplates, over time, the combination of five components. A machine-readable evidence vocabulary (present). SHACL validation shapes (present, including dual profile). Live agent-event streams emitting conformant data in production (contemplated; not yet implemented). A Standardised Supervisory Ingestion Interface (specified architecturally; not yet implemented). Runtime Autonomy Gates and a push-based Compliance MCP Server capable of authorising or blocking actions before execution (discussed as design direction; not yet implemented). When four or five of these exist and are integrated, the system shifts from documentation infrastructure to operational control infrastructure.

Five risk categories are acknowledged. *Surveillance risk*: structured evidence graphs expose agents, users, data subjects, tools, decisions, delegation chains, and policy weaknesses. *Hildebrandt-type meta-critique*: legal informational infrastructures can foreclose contestability rather than enable it; the MCS is exposed to the framing trap, the portability trap, the formalism trap, the ripple-effect trap, and the solutionism trap identified in the Hildebrandt literature. *Vendor lock-in*: commercial appropriation through restrictive derivatives or proprietary extensions of the CC BY 4.0 base. *Regulatory capture*: standards-body-level dynamics in the transition of any successor to W3C Recommendation or CEN-

CENELEC harmonised-standard status. *Weaponisation by authoritarian deployers*: a supervisor-ingestible evidence pipeline is also a surveillance pipeline; MCS adoption outside the rule-of-law context the EU assumes could produce outcomes the designers would not endorse.

The honest position on the meta-critique is that the MCS sits on a continuum, not on a side. With the three implications of Section 7.6 honoured, plus three additional safeguards not yet encoded in the MCS (a translation layer for non-technical data subjects, schema versioning with explicit recording of unrepresented attributes, periodic schema re-justification against the case-law base), the MCS occupies a position substantially toward the “through design” pole. A reviewer who treats the by/through distinction as a hard binary and the meta-critique as decisive will find the MCS wanting. A reviewer who accepts that the binary is in practice a continuum and that the honest comparator is “what supervisors face today without any compliance ontology” will find the MCS a meaningful improvement.

The split-publication strategy adopted for the release artefacts is a partial mitigation. Data-model artefacts are open. Operational-layer artefacts (SSII reference implementation, Compliance MCP Server, live connectors) remain closed until independent threat-model review and supervisor-facing validation. This removes the surface that weaponisation concerns would attack at a future date. The split is a publication-ethics decision, not a technical control: a motivated actor could build an operational layer from the open artefacts. The split buys time rather than eliminating the risk.

## 7.8 Distributive effects

If the MCS or a comparable specification became the de facto compliance norm under the AI Act and GDPR Article 22, who gains and who loses. The answer is not symmetric, and the asymmetry is not the kind that further specification work can correct.

Three classes of actor are advantaged. *Large controllers with established compliance budgets*: a controller already running ISO/IEC 27001, ISAE 3402 Type 2, and an internal model-validation function has the institutional infrastructure to absorb MCS adoption as an incremental extension. Marginal-cost asymmetry of this shape is the standard mechanism by which compliance regimes produce winner-take-most concentration. *Standards-literate consultancies and integrators*: firms that can read SHACL, deploy DPV and GDPRov, design SSII endpoints, and translate supervisor requirements into Turtle become the necessary intermediaries. Apparens is one such firm; Big 4 firms with AI assurance practices are the larger such firms. The author’s commercial position therefore aligns directly with the rise of MCS-style compliance norms; the Section 1 disclosure applies with full force. *DPA’s and supervisors with technical staff*: supervisors without such staff are operationally disadvantaged.

Three classes of actor are disadvantaged. *Small and medium-sized controllers without compliance budgets*: SMEs subject to the AI Act high-risk regime face a binary choice between hiring a consultancy at €25,000 to €100,000 per year (Apparens-tier) or a Big 4 firm at €500,000 to €2,000,000 per year, or attempting MCS adoption in-house with significant risk of error. The European Commission SME definition (fewer than 250 employees, turnover ≤ €50m or balance sheet ≤ €43m) covers a substantial share of providers of high-risk AI systems in healthcare AI, regtech, edtech, and HR-tech. The MCS does not produce this disadvantage; the AI Act and GDPR produce it. The MCS amplifies it by raising the floor of what counts as adequate compliance evidence. *Data subjects without technical representation*: a data subject who wants to contest a SHACL-validated decision needs either technical literacy (rare among the general population) or representation by an organisation with technical literacy (NGO sector concentrated in a small number of well-funded bodies: noyb, AlgorithmWatch, Statewatch, Bits of Freedom). The specification does not foreclose the prose argument; it does not enable the technical argument either. *SMEs in regulated sectors who are simultaneously controllers and data subjects of upstream agent decisions*: a small healthcare provider receiving automated underwriting from an insurer; a small retailer receiving automated risk decisions from a payment processor; a small employer receiving automated screening from an HR-tech vendor. All three absorb the worst of the first two disadvantages.

Net effect. A specification of the MCS type, adopted as a de facto compliance norm, produces a landscape in which large controllers and well-resourced supervisors interact through specialised consultancies, with SMEs and data subjects participating as second-class users. This is approximately the existing distribution of compliance literacy in the EU; the MCS does not invent it. What the MCS does is institutionalise it in a more durable form, because once an evidentiary specification is embedded in supervisory practice, the asymmetry it produces is harder to reverse than the asymmetry produced by ad-hoc compliance prose.

Three mitigations are available without abandoning the specification approach: mandatory plain-language summaries of any SHACL ValidationReport, attached as prov:wasDerivedFrom to the report itself (prototype released as `report_translator_v0_2.py`, §E.7); a reference open-source SSII implementation maintained by a non-commercial body (W3C DPVCG, CEN-CENELEC JTC 21, or an academic consortium); an SME proportionality profile within the AI Act high-risk regime that substitutes structured prose compliance narratives for SHACL-checked structural evidence on an eligibility- and scope-bounded basis (specified in §5.6 and released as `mcs_sme_profile_v0_2.ttl`). The author has specified the first and third of these; the second remains a research agenda item. A specification author who benefits commercially from the rise of the specification class is structurally limited in their capacity to design mitigations against that benefit. The specification author can identify the problem and offer a structural starting point; the specification author should not be trusted to finalise it. Follow-on work led by a different team, or funded by a body whose mission is the disadvantaged classes, is the natural locus.

---

## 8. Research agenda

---

The following sixteen items are ordered by precedence and address both theoretical foundations and practical implementation. Where empirical anchoring exists or has been ruled out, the relevant Appendix B finding is cited. The items are single-team-feasible unless otherwise flagged.

1. *EU-AIAct extension and AI Act ontology alignment.* Drive the W3C DPVCG EU-AIAct extension (v2.3, 25 February 2026) to mature status, or develop an equivalent CEN-CENELEC-anchored AI Act ontology covering Provider, Deployer, Manufacturer roles, Annex III categories, Art. 12 typed logs, Art. 72 post-market monitoring, and the NonQualifyingADM residual class. Highest-impact single action: eliminates ten of twenty-three MCS gaps simultaneously. The NonQualifyingADM class is not addressed by EU-AIAct alignment as currently scoped and would require separate ontology work.
2. *Standardised Supervisory Ingestion Interface (SSII) specification and reference endpoint.* The SSII is a transport-layer specification operating over the MCS evidence graphs. It would allow a supervisor to issue a query against a known ontology, receive machine-readable RDF, and run the same query against multiple controllers for cross-comparison. Two operating modes are specified as design direction: pull (SPARQL endpoints over MCS-conformant graphs with supervisor-side query templates for the five standard questions) and push (a Compliance MCP Server that emits structured evidence on a configured schedule or on trigger conditions). Open items: threat model covering authentication, authorisation, audit logging of supervisor queries, tenant isolation between regulated entities, per-query scope limitation, and rate limiting; minimum-viable reference endpoint over pyshacl-validated MCS graphs; at least one supervisor pilot to validate query-template adequacy. The SSII is treated here as research direction rather than delivered contribution; retention of the name across versions is a commitment to the implementation work, not a claim of current delivery. The capability this addresses (GDPR Article 60 cooperation and EDPB Article 65 binding-decision function operating at supervisory-ingestion volume) is operationally needed but cannot today be exercised at anything approaching the volume of cases supervisors would need.
3. *Temporal Delta Auditing system.* Technical infrastructure to reconstruct how agent reasoning would differ under alternative policy versions, implementing versioned Ontology Ledgers with semantic differencing and lex mitior compliance evaluation. The lex mitior construction is doctrinally familiar from criminal law; its application to administrative compliance has no case-law anchor as of April 2026. Frame any pilot as a doctrinal proposal.
4. *Empirical validation of the Profile A quantitative pairings.* Do the ( $\epsilon = 10\%$ ,  $M = 50\%$ ) pairings, applied at controller-side, predict the Solely versus Assisted classifications supervisors and courts subsequently endorse. Requires controller participation in disclosing flip-rate measurements alongside ex-post supervisory determinations; cross-controller and cross-sector aggregation; statistical analysis of the pairing against endorsed outcomes.
5. *Meaningful-intervention predicate.* The HumanIntervention six-property schema encodes WP251rev.01 plus AP handvatten criteria plus the authorityActuallyExercised property responding to Uber/Ola. Whether the six properties jointly predict the substantive-review test courts apply is an empirical question. Empirical anchor: Section 4.6 Finding 3 (Uber/Ola Amsterdam Court of Appeal April 2023 reversal; Krakow-office symbolic-review failure mode). Research question: interaction-log metrics that predict Solely versus Assisted sufficiently in advance of litigation for controllers to self-correct.
6. *Causal blast radius algorithms for revocation propagation.* Formal-methods-grade algorithms for computing, at revocation time, the set of agents, tasks, queued work, cached artefacts, and downstream decisions that must be acted on. The natural site for LTL or CTL verification within the MCS.
7. *Authority-utilisation metrics for human intervention.* Deviation rate, deviation latency, time-on-task, authority-exercised-frequency as predictors of the substantive-review test. Empirical.
8. *Evidence decay algorithms for autonomous systems.* Sector-specific evidence-weight decay functions with empirically validated refresh cadences. Per Section 6.2, no current empirical anchor for any specific cadence.
9. *Neurosymbolic auditability for LLM agents.* Chain-of-Thought reasoning that natively outputs RDF triples aligned with enterprise ontologies. Eliminates post-hoc extraction and enables direct supervisory ingestion of reasoning traces.
10. *Dynamic risk re-classification.* Runtime monitoring detecting when autonomous goal-setting moves an agent from Low-Risk to High-Risk Annex III categories, with automatic triggering of enhanced logging and oversight. RC-flagged conjecture.
11. *European Agent Identity Registry.* Legal framework for persistent agent identity building on SPIFFE/SPIRE standards, with traceable legal personas across platforms and jurisdictions.
12. *Resolving the Article 22 rights-vs-procedural-obligation debate ontologically.* Whether agent ontologies should encode DataSubjectInvocationEvent classes (rights framing) or continuous ComplianceAssessment classes (procedural accountability framing), or both.
13. *Bridge artefact maintenance and Member State variants.* Ongoing maintenance of `mcs_framework_bridges.ttl` to track NIS2 transposition, Italian Law 132/2025 equivalents, and sectoral-extension ontologies as they mature. Italy and UK are the first two national-variant instances committed to for the follow-up revision.
14. *Validation-pack execution and aggregation.* Operate the three open tracks of the companion

validation pack. Track 1 (inter-rater kappa on the 25-case sample; committed for the follow-up revision). Track 2 (SHACL throughput at 1, 10, 100, 1,000 decisions per second on real runtimes). Track 3 (structural fit across topologies). Roll into subsequent revisions as contributor reports arrive.

15. *Open test corpora for retrospective supervisory reconstruction*. Synthetic but realistic agent-fleet event streams with deliberately introduced policy-version transitions, delegation grants and revocations, ontology-version updates; reference reconstructions with both action-time and current-time policy evaluations explicit; labelled defect cases. Natural maintainer is an academic consortium or a CEN-CENELEC-adjacent expert group.
16. *Privacy-preserving provenance patterns*. PROV-O completeness is in tension with GDPR Art. 5(1) (c) data minimisation. Research question: the minimum personal-data retention compatible with full authority and decision-trace reconstruction at the supervisory level.

---

## 9. Conclusion

---

Agentic AI has outpaced the ontologies intended to govern it. Vendor ontologies model enterprise workflow fluently but ignore data-protection primitives. The open normative stack (PROV-O, OWL-Time, SHACL) supplies the spine but not the legal-semantic layer. GDPRov and DPV supply the GDPR layer but not the AI Act roles, the Article 22 Decision taxonomy, or the Article 12 typed log events. ISO/IEC and CEN-CENELEC supply prose, not machine-readable ontology. No instrument combines these into a supervisor-ingestible whole.

The MCS proposes such a combination. 23 conceptual slots organised into a twelve-slot agent-behaviour core and an eleven-slot supervisory-evidence layer, anchored in operative statutory provisions or in empirical case-law validation, contributing 16 net-new classes under strict reuse-zero accounting (69 owl:Class declarations in the source Turtle file). A temporal and validation layer expressed in OWL-Time with fourteen core SHACL shapes plus seven profile-layer shapes in v1.2 plus a Bayesian evidence-decay structure with declared sectoral priors. An integration layer reusing GDPRov, DPV, and PROV-O via owl:imports. A Standardised Supervisory Ingestion Interface identified as research direction in Section 8 rather than claimed as delivered contribution. A companion validation pack with three pre-registered open tracks, of which Track 1 is committed for the follow-up revision.

The central claim is precise and bounded. The MCS is a working paper, not a reference architecture. The 23-item claim refers to conceptual slots; the 16-item figure is the net-new-class count under strict reuse-zero accounting; the 69-item figure is the raw owl:Class count in the Turtle source. No competing synthesis the author has reviewed achieves the combination of GDPR legal-semantic anchoring via DPV/GDPROV, three-way Decision taxonomy with dual-operationalisation drawsStronglyOn predicate, dual-profile SHACL for supervisory divergence with Profile A (quantitative, AP-inspired permissive) and Profile B (qualitative, CNIL/German-guidance-inspired stricter) both shipped in v1.2, size-based SME proportionality profile as a separate axis, Bayesian evidence-decay with declared sectoral priors, AI Act typed log events at statutory-subclass granularity, and the NonQualifyingADM residual class grounded in a 25-case empirical sample. Cobbe et al. (2025) is closest on supervisor-ingestibility but does not commit to specific ontological structures. Agent Spec is closest on portability but fragments audit trails. ACP Framework is closest on delegation semantics but lacks GDPR legal primitives. If a reader knows of a synthesis that achieves the combination with fewer classes, closer statutory fit, or better empirical anchoring against the same 25-case sample, the author commits to engaging it in a revised edition.

The claim has genuine limitations. Empirical validation is internal-consistency only; Track 1 is committed for a follow-up revision. Structural completeness is not legal compliance, and the SHACL structural-conformance flag cannot substitute for supervisory judgement. The SSII is research direction, not delivered; its promotion to Section 8 rather than retention in the main design is honest positioning. The specification remains under Apparens commercial authorship; standards-body governance (W3C DPVCG or CEN-CENELEC JTC 21) is Section 8 item 1. The treatment of multi-stakeholder supply chains is architectural rather than interface-specified. The distributive effects of adoption as a de facto compliance norm are regressive; the size-based SME proportionality profile (§5.6) is a partial structural response but cannot correct the underlying asymmetry. The external regulatory coherence on which the MCS depends is empirically contingent, with NIS2 transposition incomplete and Italian Law 132/2025 signalling heterogeneity at the national level.

These limitations are acknowledged without apology. Every governance framework in the reviewed corpus shares the empirical validation deficit. The technical enforcement gap is by design: the MCS governs at the evidence layer, not the enforcement layer. The distributive-effects limitation is genuine and cannot be repaired by the author. The other limitations are actionable and are in the research agenda.

The MCS succeeds if European supervisors, regulated controllers, and peer researchers find the synthesis defensible enough to build on and sharp enough to critique. It fails if that community finds the synthesis either non-reproducible (coders disagree under the same rubric), operationally infeasible (SHACL does not run at production throughput), or structurally inadequate (real agent fleets require more than three extensions to be represented). The validation pack's three tracks test these three failure modes in order. Both outcomes, success and failure, are scientifically useful and both are publishable as version 2. That commitment is the paper's final claim.

---

## Data and code availability

---

This paper and its companion artefacts are deposited on Zenodo under CC BY 4.0:

- **Version DOI (v0.5.1):** [10.5281/zenodo.19758441](https://doi.org/10.5281/zenodo.19758441)
- **Concept DOI:** the version DOI's parent record on Zenodo resolves to the latest release.

The deposit contains: - The paper PDF (`mcs_paper_v0_5_1.pdf`). - The framework specification (`mcs_framework_specification_v0_5_1.pdf`). - The companion pack PDF bundling paper, framework spec, and 12 source-file appendices (`mcs_companion_pack_v0_5_1.pdf`). - The complete source pack as a zip (`MCS_v0_5_1_complete_pack.zip`) containing the OWL/RDF vocabulary (`mcs_vocabulary.ttl`), SHACL profiles (`mcs_profiles_v1_2.ttl`, `mcs_sme_profile_v0_2.ttl`), TLA+ revocation model (`mcs_revocation_model_v0_2.tla` plus `.cfg`), case sample (`mcs_case_sample_v0_3.csv`), mapping matrix (`mcs_mapping_matrix_v0_3.csv`), deep-coding matrix, scenarios, validation pack, reproducibility scripts, SHACL examples, Track 1 blinded coder package, Track 2 synthetic benchmark, and the legal mapping review template.

Reproducibility entry point: `reproduce_package_counts_v0_5_1.py`. Verification battery: see `ERRATA_v0_5_1.md` and `mcs_artifact_manifest_v0_5_1.md` in the deposit. Pending external evidence is itemised in `IDEAL_CORRECTION_STATUS_v0_5.md` and `tla_model_checking/TLC_STATUS_v0_5.md`.

---

## References

---

- Argyris, C. and Schön, D. A. (1978) *Organizational Learning: A Theory of Action Perspective*. Reading, MA: Addison-Wesley.
- Atkinson, K., Bench-Capon, T. and Bollegala, D. (2020) Explanation in AI and law: Past, present and future. *Artificial Intelligence*, 289.
- Autoriteit Persoonsgegevens (2024) *Advies artikel 22 AVG en geautomatiseerde selectie-instrumenten*, 10 October 2024.
- Autoriteit Persoonsgegevens (2025) *Handvatten betekenisvolle menselijke tussenkomst*, July 2025.
- Autoriteit Persoonsgegevens and Rijksinspectie Digitale Infrastructuur (2024) *Final advice on the supervisory structure for the AI Act*, 7 November 2024.
- AFM (2026) *AFM Agenda 2026*, 19 January 2026.
- Bench-Capon, T. and Sartor, G. (2003) A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence*, 150(1-2), pp. 97-143.
- BaFin (2025) Guidance on ICT risks in the use of AI at financial entities, 18 December 2025.
- CJEU (2023) Case C-634/21 *SCHUFA Holding (Scoring)*, judgment of 7 December 2023.
- CJEU (2023) Case C-26/22 *SCHUFA Holding (Discharge)*, judgment of 7 December 2023.
- CJEU (2025) Case C-203/22 *CK v Dun and Bradstreet Austria*, judgment of 27 February 2025.
- Cobbe, J., Singh, J. and Sheridan, T. (2025) Governance and accountability frameworks for AI agents. Working paper.
- DNB (2019) *General principles for the use of Artificial Intelligence in the financial sector (SAFEST)*, 25 July 2019.
- European Commission (2016) Regulation (EU) 2016/679 — General Data Protection Regulation. *OJ L* 119, 4 May 2016.
- European Commission (2022) Directive (EU) 2022/2555 — NIS2 Directive. *OJ L* 333, 27 December 2022.
- European Commission (2024) Regulation (EU) 2024/1689 — Artificial Intelligence Act. *OJ L* series, 12 July 2024.
- European Commission (2024) Directive (EU) 2024/2853 — Revised Product Liability Directive. *OJ L* series, 18 November 2024.
- Hashmi, M., Governatori, G., Lam, H.-P. and Wynn, M. T. (2018) Are we done with business process compliance: state of the art and challenges ahead. *Knowledge and Information Systems*, 57, pp. 79-133.
- Hildebrandt, M. (2015) *Smart Technologies and the End(s) of Law*. Cheltenham: Edward Elgar.
- Janssen, J. (2025) *The AI Accountability Trap*. Deventer: Apparens.
- Janssen, J. (2026) From Battlefield to Boardroom: Strategic Red Teaming as an Epistemic Governance Instrument in the Age of AI. Apparens Working Paper.
- Janssen, J. (2026) The Implementation Gap: The AI Enterprise Control Index as an Operational Governance Instrument for Agentic AI Systems. Apparens Working Paper.
- Klein, G. A. (1993) A recognition-primed decision (RPD) model of rapid decision making. In Klein, G. A. et al. (eds.) *Decision Making in Action*. Norwood, NJ: Ablex.
- Knublauch, H. and Kontokostas, D. (2017) Shapes Constraint Language (SHACL). W3C Recommendation.

- Lebo, T., Sahoo, S. and McGuinness, D. (2013) PROV-O: The PROV Ontology. W3C Recommendation.
- Mittelstadt, B. (2019) Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence*, 1(11), pp. 501-507.
- Pagallo, U. (2013) *The Laws of Robots: Crimes, Contracts, and Torts*. Dordrecht: Springer.
- Pandit, H. J. and Lewis, D. (2017) Modelling provenance for GDPR compliance using linked open data vocabularies. SEMANTICS 2017.
- Pandit, H. J., et al. (2019) Creating a vocabulary for data privacy (DPV). *OTM Confederated International Conferences*.
- Sartor, G. (2005) *Legal Reasoning: A Cognitive Approach to the Law*. Dordrecht: Springer.
- Sergot, M. J., Sadri, F., Kowalski, R. A., Kriwaczek, F., Hammond, P. and Cory, H. T. (1986) The British Nationality Act as a logic program. *Communications of the ACM*, 29(5), pp. 370-386.
- Veale, M. and Edwards, L. (2018) Clarity, surprises, and further questions in the Article 29 Working Party draft guidance on automated decision-making and profiling. *Computer Law and Security Review*, 34(2).
- Verheij, B. (2003) Dialectical argumentation with argumentation schemes: An approach to legal logic. *Artificial Intelligence and Law*, 11(2-3).
- Verheij, B. (2017) Formalizing arguments, rules and cases. *ICAIL 2017*.
- Wachter, S., Mittelstadt, B. and Floridi, L. (2017) Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. *International Data Privacy Law*, 7(2), pp. 76-99.
- Article 29 Data Protection Working Party (2018) Guidelines on automated individual decision-making and profiling (WP251rev.01).
-

# Appendix A: Operational mapping and formal sketches

---

## A.1 Matrix structure

---

The mapping matrix rates eight agent-behaviour primitives (rows) against six legal categories (columns). Each cell carries a status flag (SD statutory derivation, IP interpretive proposal, RC research conjecture), a confidence grade (H high, M medium, L low), and a pointer to the formal sketch in A.4 through A.19 that derives the rule where one exists.

## A.2 Legend

---

Primitives (rows): Tool invocation, Memory read, Memory write, Delegation grant, Delegation revocation, Policy version change, Autonomous goal setting, Cumulative goal execution.

Legal categories (columns): GDPR Art. 22, GDPR Art. 28, AI Act Art. 12, AI Act Art. 14, AI Act Annex III, NIS2 Art. 23.

## A.3 Table A.1 (Profile A coding)

---

The full 8 by 6 matrix under Profile A (AP-inspired permissive). Representative cells plus NIS2-parity sketches are reproduced in compressed form here; the complete matrix with per-cell justifications is maintained in `mcs_mapping_matrix_v0_3.csv` in the release repository.

Sources per column. GDPR Art. 22: SCHUFA paras 50, 73; Dun and Bradstreet paras 40, 59 to 60; WP251rev.01; AP handvatten July 2025. GDPR Art. 28: Art. 28(2), 28(3), 28(4). AI Act Art. 12: Art. 12(2) (a) to (c); Art. 12(3). AI Act Art. 14: Art. 14(1), 14(4)(a) to (e); AP-RDI Final advice 7 November 2024. AI Act Annex III: Art. 6(2), 6(3); Annex III categories 1 to 8. NIS2 Art. 23: Art. 23(1), 23(3), 23(4); Recital 101.

Cell-level notes follow the formal sketches in A.4 onwards. The NIS2 column carries four F-sketches (F4, F11, F12, F13) for parity with the GDPR column.

## A.4 F1 Decision × Art. 22

---

See Section 5.3 inline sketch.

## A.5 F2 Score × SCHUFA upstream

---

See Section 5.3 inline sketch. Operationalisation splits per profile.

## A.6 F3 Tool invocation × effect propagation

---

**Status:** interpretive proposal. **Source:** Art. 22 via tool-chain; no direct authority.

```
∀a, t, d. Agent(a) ∧ ToolInvocation(t) ∧ Decision(d) ∧
  a.performed(t) ∧ t.contributedTo(d) →
  AttributableToAgent(d, a)
  ∧ ( drawsStronglyOn(d, t.output) → ToolOutputInScope_Art22(t) )
```

Captures tool-chain propagation of Article 22 engagement through agent action. Does not capture edge cases where a tool invocation produces an effect but the effect propagates through a separate decision actor.

## A.7 F4 Tool invocation × NIS2 significant incident

---

See Section 5.3 inline sketch. Heaviest-load NIS2 cell.

## A.8 F5 Memory read × Art. 22 contributing factor

---

**Status:** research conjecture. **Source:** none.

```
∀m, d. MemoryRead(m) ∧ Decision(d) ∧ m.precedes(d) ∧
  m.readValue.contributedTo(d) →
  Art22InputEvent(m)
  (conjecture; no statutory anchor)
```

## A.9 F6 Delegation × Art. 28 processor chain

---

**Status:** statutory derivation. **Source:** GDPR Art. 28(2), (4); Art. 28(3).

```
∀d1, d2. DelegationGrant(d1) ∧ DelegationGrant(d2) ∧
  d2.isSubdelegationOf(d1) →
  ( d1.hasScopeConsistentWith(d2.scope) )
  ∧ ( d1.hasWrittenAuthorisation ∨ d1.hasRecordedAuthorisation )
  ∧ ( Revoked(d1) → Revoked_propagatesTo(d2) )
```

Captures the Art. 28(2) prior-authorisation requirement; the Art. 28(3) scope-consistency requirement;

transitive revocation propagation through delegation chains.

## A.10 F7 Policy version × lex mitior

**Status:** research conjecture. **Source:** none directly.

```
∀d, p1, p2. Decision(d) ∧ Policy(p1) ∧ Policy(p2) ∧
d.governedBy(p1) ∧ p2.supersedesPolicy(p1) ∧
p2.isMorePermissive(p1) →
lex_mitior_available(d, p2) (conjecture)
```

## A.11 F8 Evidence × AS 1105

**Status:** interpretive proposal. **Source:** PCAOB AS 1105 reliability hierarchy.

```
∀e. EvidenceArtifact(e) →
ReliabilityWeight(e.source) ∈
{ SystemLogSource ↦ 0.85,
  HumanAttestationSource ↦ 0.60,
  ExternalReportSource ↦ 0.70,
  ThirdPartyAuditSource ↦ 0.90 }
(illustrative defaults; controller-configurable)
```

## A.12 F9 Human intervention × authority utilisation

**Status:** statutory derivation (predicates); interpretive proposal (thresholds). **Source:** WP251rev.01; AP handvatten July 2025; Uber/Ola Hof Amsterdam.

```
∀hi. HumanIntervention(hi) →
MeaningfulIntervention(hi) ↔
( hi.hasAuthorityToDeviate = true
  ∧ hi.hasUnderstandingOfLogic = true
  ∧ hi.hasMarginOfDiscretion = true
  ∧ hi.consideredAllRelevantData = true
  ∧ hi.authorityActuallyExercised = true )

AuthorityActuallyUtilised(hi) ↔
population-level deviation rate > threshold θ
(illustrative θ = 10%; not anchored in case-law)
```

## A.13 F10 Risk classification × Annex III

**Status:** statutory derivation. **Source:** AI Act Art. 6(2), 6(3), Annex III.

```
∀s. AISystem(s) →
HighRisk(s) ↔
( IntendedForAnnexIIIUseCase(s)
  ∧ ¬(s matches Art. 6(3) derogation criteria)
  ∧ ¬(s matches Art. 6(3a) re-classification criteria) )
```

## A.14 F11 Delegation revocation × NIS2 incident window

**Status:** interpretive proposal. **Source:** NIS2 Art. 23(3)-(4); EDPB Guidelines on notification of personal data breaches (2017) for temporal-window reasoning pattern.

```
∀dr, i. DelegationRevocation(dr) ∧ Incident(i) ∧
SignificantIncident(i) ∧ dr.occursWithin(i.detectionWindow) →
( dr MUST be preserved in the incident evidence set
  ∧ dr.propagationEffects MUST be recorded as part of i.containment
  ∧ Revoked(dr.parent) → StopsInflightWork applies by default
  DURING incident window rather than default BlocksNewWork+MarksArtifactsStale )
```

Captures three things. First, NIS2 Art. 23 notification obligations extend to the set of delegation events within the incident's temporal window; evidence preservation under Art. 23(4) includes revocation events, not only positive actions. Second, revocation propagation effects are part of incident containment under Art. 23(4)(c) and must be auditable. Third, default revocation effects are sector-calibrated: during an incident window, StopsInflightWork becomes a defensible default where ordinarily it is opt-in (Section 6.1). Does not capture Member State transposition variants where the 24h/72h reporting windows of Art. 23(4) differ.

## A.15 F12 Memory write × NIS2 Art. 23(4) reporting cadence

**Status:** interpretive proposal. **Source:** NIS2 Art. 23(4); Recital 102.

```
∀mw, i. MemoryWrite(mw) ∧ Incident(i) ∧ SignificantIncident(i) ∧
mw.writtenDuring(i.window) →
( mw MUST be preserved under WORM semantics from i.detectionTime
  until i.finalReport.filed + Art23(6).reviewPeriod
  ∧ mw.metadata MUST include source agent, delegation grant id,
  and policy version active at write time )

EarlyWarning(i) → 24h deadline from detection
IncidentNotification(i) → 72h deadline from detection
FinalReport(i) → 1 month after notification
```

Captures the NIS2 staged reporting cadence as applied to agent memory writes during an incident window: memory writes that could bear on the incident must be preserved across the cadence, with sufficient metadata to reconstruct the agent-delegation-policy triple active at write time. The stringent preservation obligation derives from Art. 23(4) evidentiary requirements plus Art. 23(6) follow-up review. Does not capture cross-border incidents under Art. 23(8) where the deadlines may be adjusted by CSIRT cooperation arrangements.

### A.16 F13 Policy version change × NIS2 Art. 21 measures

---

**Status:** interpretive proposal. **Source:** NIS2 Art. 21 cybersecurity risk-management measures; Annex I.

```
∃pv1, pv2, e. Policy(pv1) ∧ Policy(pv2) ∧ pv2.supersedes(pv1) ∧
  Entity(e) ∧ e.scope ∈ {Essential, Important} ∧
  pv1.isNIS2MeasureAttestation ∨ pv2.isNIS2MeasureAttestation →
  ( pv1 → pv2 transition MUST be recorded with
    Art21(2)-mapped measure deltas
  ∧ IfMeasureDowngraded(pv1, pv2) →
    requires documented risk-acceptance by entity's management body
    per Art. 20(1) )
```

Captures two things. First, any policy version change that affects a NIS2 Art. 21 measure attestation is itself a governance event: the delta between pv1 and pv2 must be recordable at Art. 21(2) measure granularity (risk analysis, incident handling, business continuity, supply chain security, etc.). Second, measure downgrades trigger the Art. 20(1) management-body approval requirement; the MCS must be able to represent that approval as a first-class evidence artefact linked to the policy version transition. Does not capture the substantive standard for when a measure counts as downgraded.

### A.17 F14 Autonomous goal setting × Annex III re-classification (existing, retained)

---

**Status:** research conjecture. **Source:** AI Act Art. 6(2)-(3); no direct jurisprudence.

See `mcs_mapping_matrix_v0_3.csv` for full derivation.

### A.18 F15 Cumulative goal execution × Art. 22 (existing, retained)

---

**Status:** research conjecture. **Source:** none directly.

### A.19 F16 Delegation × AI Act Art. 14 human oversight (existing, retained)

---

**Status:** interpretive proposal. **Source:** AI Act Art. 14(1), 14(4).

---

## Appendix B: 25-case empirical sample

### B.1 Purpose

The 25-case sample is the empirical validation base for the MCS. Section 4.6 promotes the four key findings into the paper body; this appendix documents the case-level record, sources, and classifications that ground those findings.

### B.2 Selection criteria

See Section 2.3 (Criterion A and Criterion B).

### B.3 Coder and coding date

Round 1: 23 April 2026, single coder (author). Round 2 self-consistency: 7 May 2026, intra-rater. Track 1 independent coder: committed for a follow-up revision (Section 8 item 14).

### B.4 Case table

Table B.4. Twenty-five EU ADM enforcement cases (25 May 2018 to April 2026).

#	Jurisdiction	Date	Brief facts	Classification	Source
B1	CJEU, Grand Chamber (C-634/21 <i>SCHUFA</i> )	7 Dec 2023	Credit-reference agency's automated probability score; bank refused loan; CJEU held score itself a "decision" under Art. 22(1) where downstream draws strongly on it.	Solely	Judgment paras 50, 73
B2	CJEU (C-203/22 <i>Dun and Bradstreet Austria</i> )	27 Feb 2025	Mobile-phone operator denied contract based on D&B automated creditworthiness assessment; "lacked any manual oversight".	Solely	Judgment paras 40, 59-60
B3	Verwaltungsgericht Wiesbaden	Jan 2026	SCHUFA TeamBank specialised scorecard rated plaintiff at 85.96%; bank refused EasyCredit application; court applied SCHUFA test qualitatively.	Solely	Wiesbaden ruling
B4	Italian Garante <i>Mevaluate Onlus</i>	May 2021 (Cassazione confirm)	Platform assigned alphanumeric reputational rating.	Solely (Art. 22(2)(c) consent held invalid)	FPF Case 5
B5	Italian Garante <i>Deliveroo "Frank"</i>	July 2021	Gig-worker shift allocation and penalty algorithm.	Solely	FPF Case 3
B6	Italian Garante <i>Foodinho</i>	July 2021	Rider slot allocation.	Solely	FPF Case 6
B7	Italian Garante / Cassazione	Jan 2018 / Nov 2021	Car-sharing personalised rates.	Solely (borderline)	FPF Case 2
B8	Belgian APD	2019 (opinion)	Automated authorisation of collective self-consumption by Wallonia energy authority.	Solely (Art. 22(2)(b))	FPF §1.4
B9	Amsterdam District Court <i>Ola</i>	Mar 2021	Driver allocation, fraud-probability scores.	Solely	11 March 2021
B10	Amsterdam Court of Appeal <i>Uber</i> robo-firing	4 Apr 2023	Four drivers dismissed after fraud allegation;	Solely (REVERSED first instance)	ECLI:NL:GHAMS:2023:793

			Krakow review “not much more than a purely symbolic act”.	Assisted)	
B11	Amsterdam Court of Appeal <i>Uber</i> data access	4 Apr 2023	Six drivers’ Art. 15 requests; disclosure of profiling logic ordered.	Solely	Hof Amsterdam 4 April 2023
B12	Amsterdam Court of Appeal <i>Ola</i>	4 Apr 2023	Driver requests on fraud-probability and earning profile.	Solely	Hof Amsterdam 4 April 2023
B13	Court of First Instance The Hague	Feb 2020	Gun-licence e-screener.	Assisted	FPF Case 8
B14	Austrian Federal Admin Court AMS AMAS	Dec 2020	Employability algorithm “to assist counsellors”.	Assisted	FPF Case 9
B15	Slovak Constitutional Court <i>e-kasa</i>	10 Nov 2021	Real-time receipt data feeding automated risk profiles.	Assisted (court declined Art. 22)	FPF Case 4
B16	Amsterdam District Court <i>Uber</i> first instance	11 Mar 2021	Four drivers’ deactivation; “whole team” decided.	Assisted at first instance; Solely on appeal	Rb Amsterdam 11 March 2021
B17	Norwegian Datatilsynet <i>IB grading</i>	Aug 2020	COVID-era automated IB grading; decided on Art. 5/25.	Out-of-Art.22-Scope	FPF Case 7
B18	HmbBFDI / ICO / CNIL / Garante <i>Clearview AI</i>	2021-2022	Facial recognition database from web-scraped images; decided on Art. 6/9.	Out-of-Art.22-Scope	FPF Cases 10-13
B19	CNIL / Datainspektionen live FR in schools	2019-2020	Facial recognition for school access; decided on Art. 6/9.	Out-of-Art.22-Scope	FPF Cases 30, 31
B20	AP <i>Belastingdienst Toeslagen</i>	7 Dec 2021	€2.75M fine; discriminatory processing of nationality data. Decided on Art. 5(1)(a)/6.	Out-of-Art.22-Scope	AP boetebesluit 7 Dec 2021
B21	AP <i>Belastingdienst toezichtarrangement</i>	Aug 2025	AP requested plan-of-action on selectie-instrumenten.	Pending	AP letter 8 Aug 2025
B22	District Court The Hague <i>SyRI</i>	5 Feb 2020	Benefits-fraud algorithm; ruled unlawful under Art. 8 ECHR.	Out-of-Art.22-Scope	ECLI:NL:RBDHA:2020:865
B23	Italian Cassazione <i>Mevaluate</i> appeal	25 May 2021	Confirmed Garante.	Solely	Cassazione 25 May 2021
B24	CNIL <i>Clearview</i> fine	Oct 2022	€20M fine; decided on Art. 6/9.	Out-of-Art.22-Scope	CNIL Oct 2022
B25	Hamburg DPA <i>SCHUFA</i> (post-judgment)	Mar 2024	DPA position on SCHUFA extension to AI credit scoring.	Interpretive (Solely default)	Hamburg DPA position

## B.5 Findings

See Section 4.6 (promoted). Full per-case rationales and borderline-case sensitivity analysis retained in `mcs_case_sample_v0_3.csv` in the release repository.

## B.6 Qualitative vs quantitative drawsStronglyOn

See Section 6.4.

## **Appendix C: Worked scenarios (moved to companion)**

---

The seven worked-example scenarios (retail banking, hospital triage, municipal benefits, platform delivery, revocation propagation, cross-border one-stop-shop, NIS2 plus GDPR concurrent incident) are maintained in the companion scenarios document `mcs_scenarios_v0_2.md` to keep the main paper focused on the core contribution. Scenario summaries are referenced in Section 7.1.

---

## Appendix D: Validation pack architecture

---

### D.1 Three-track design

---

The validation pack is published as a companion document (`validation_pack_v1.md`). It specifies three pre-registered open tracks.

**Track 1: Independent Replication.** One independent coder applies the Appendix E.6 coding rubric to the 25-case Appendix B sample, blind to the author's round-1 classifications. Cohen's kappa computed per artefact. Pre-registered hypothesis:  $\kappa \geq 0.70$ . Failure criterion:  $\kappa < 0.60$  triggers rubric revision. **Committed for the follow-up revision.**

**Track 2: Technical Feasibility Benchmark.** Contributor emits MCS-conformant Turtle at 1, 10, 100, 1000 decisions per second; measures SHACL validation latency on pyshacl, Apache Jena SHACL, TopBraid SHACL API. Pre-registered hypothesis: p95 validation latency  $< 100$ ms at 100 decisions per second.

**Track 3: Structural Fit Pilot.** Contributor with a different agent topology validates whether every event class can be represented by an MCS class without requiring extension. Pre-registered hypothesis: structural sufficiency.

### D.2 Eligibility for contributor pilots

---

The validation pack is open to any contributor meeting the per-track eligibility criteria specified in `validation_pack_v1.md`. No specific pilots are pre-identified; submissions are evaluated on a rolling basis as they arrive. The author commits to acknowledging receipt of any qualifying submission and reporting outcomes regardless of the contributor identity or affiliation.

### D.3 Open-call mechanics

---

See `validation_pack_v1.md` for full mechanics. Eligibility: working AI agent system operating at least six months; technical capacity to populate Turtle and run SHACL. Submission window 12 weeks from pack receipt with one extension permitted.

### D.4 What the validation pack does and does not produce

---

Produces: independent reliability evidence; technical feasibility measurements; structural-fit findings.

Does not produce: regulatory validation (no DPA has endorsed the MCS; no CJEU has applied it; no notified body has cross-mapped it). Regulatory validation is a separate programme.

---

## Appendix E: Coding rubrics for reproducibility

---

### E.1 Purpose and scope

---

Five coded artefacts in the paper. This appendix publishes the rubrics. Full protocols with CSV schemas, seeded shuffle, tie-breakers, randomisation, and materials access maintained in `coding_rubrics_v0_2.md` in the release repository.

### E.2 Survey coverage matrix prompt

---

Three-value rating {present, partial, absent} per cell (concept × ontology). Present: formal construct matches concept under synonymous or strictly narrower terms. Partial: prose-level coverage without formal construct, or identifiable-but-imperfect overlap. Absent: no overlap. See `coding_rubrics_v0_2.md` §2 for full step rules.

### E.3 Deep-coding reusability matrix prompt

---

Three-value rating {reusable, partially reusable, not reusable} per cell (MCS class × ontology). Reusable requires operational-criteria preservation under owl:imports. See `coding_rubrics_v0_2.md` §3.

### E.4 Mapping matrix status and confidence prompt

---

Status {SD, IP, RC} plus confidence {H, M, L} per cell (primitive × legal category). Profile A and Profile B codings separate. See `coding_rubrics_v0_2.md` §4-5.

### E.5 Case classification prompt

---

See Section 2.7 (promoted to main body). Full five-step protocol with sector-calibrated thresholds in `coding_rubrics_v0_2.md` §6.

### E.6 Limits

---

Prompt publication does not eliminate interpretive judgement; operationalised thresholds are defensible but not statutory; SD/IP boundary is itself interpretive; language-tier coding does not fully correct for selection bias; COI disclosures are self-reported; round-2 audit is intra-rater not inter-rater.

### E.7 Plain-language ValidationReport translator

---

A SHACL `sh:ValidationReport` is machine-readable RDF with result blocks identifying which shapes failed, what targets they applied to, and what severity the failure carried. It is not readable by a non-technical data subject, an SME controller without SHACL literacy, or a supervisory caseworker without RDF training. The distributive-effects analysis in §7.8 identifies this as a specific second-order exclusion.

The MCS releases `report_translator_v0_2.py` as a reference implementation (~300 lines, Python 3.11, depends on `rdflib` and `pyshacl`). The translator takes a `sh:ValidationReport` and produces a plain-language narrative attached to the report itself as `prov:wasDerivedFrom`. The narrative has four sections: what was validated (target class and instance count), what passed (shape names and their natural-language equivalents), what failed and why (shape, target, violation message, severity), and what this does and does not mean (the §6.3 and §7.3 disclaimers that structural completeness is not legal compliance, rendered in non-technical language).

The translator is reference implementation, not production code. It does not attempt to translate arbitrary SHACL severity chains, it does not render profile-delta governance events beyond naming the profiles involved, and it has no authentication or multi-language support in this release. The intent is to demonstrate that the translation is mechanically feasible so that downstream controllers, SMEs, NGOs, and data-subject representatives can build production variants for their specific audience.

Licence: CC BY 4.0 (more permissive than the paper's CC BY-NC-ND to allow downstream modification and commercial reuse by data-subject representation organisations). The translator's licence is deliberately distinct from the paper's licence on the grounds that mitigation infrastructure should have fewer friction points than specification text.

---

# Appendix F: Revocation propagation, LTL formalisation

## F.1 Purpose

§6.1 specifies five effects that a delegation revocation may produce (`BlocksNewWork`, `CancelsQueuedWork`, `StopsInflightWork`, `RequiresReview`, `MarksArtifactsStale`). §8 research-agenda item 6 identifies revocation propagation as the natural site for formal verification within the MCS. This appendix publishes a Linear Temporal Logic (LTL) formalisation of three of those effects plus a transitive-revocation property, at sufficient depth to show what formal verification would establish and what it would not. The formalisation is not a full verification artefact; a full artefact would include the transition system specification, the model-checker input files, and a verification transcript. This appendix publishes the formulas; the transition system lives in `mcs_revocation_model_v0_2.tla` in the release repository (TLA+ specification, not yet model-checked against a concrete model).

## F.2 Vocabulary

Atomic propositions parameterised by delegation grants  $d$ , work units  $t$ , agents  $a$ :

- $R(d)$ : delegation  $d$  is in a revoked state
- $A(a, d)$ : agent  $a$  is currently authorised under delegation  $d$
- $W(t, d)$ : work unit  $t$  is executing under authority of delegation  $d$
- $Q(t, d)$ : work unit  $t$  is queued (not yet executing) under  $d$
- $New(t, d)$ : work unit  $t$  was accepted into the work queue under  $d$  in this time step
- $Sub(d1, d2)$ :  $d2$  is a direct sub-delegation of  $d1$
- $Sub^*(d1, d2)$ :  $d2$  is in the transitive closure of  $Sub$  under  $d1$  (reflexive-transitive)
- $Stale(x, d)$ : artefact  $x$  is marked stale as a consequence of  $d$ 's revocation
- $InFlight(t)$ : work unit  $t$  is executing (not queued, not completed)
- $Opt(StopsInflight, d)$ : the `StopsInflightWork` opt-in applies to revocation of  $d$

LTL operators:  $\square$  (always),  $\diamond$  (eventually),  $\circ$  (next),  $U$  (until).

## F.3 Formulas

**F.3.1 Transitive revocation.** Revoking a delegation eventually revokes all its sub-delegations.

$$\square (R(d1) \wedge Sub^*(d1, d2) \rightarrow \diamond R(d2))$$

The formula admits unbounded latency in propagation. A bounded variant, which a supervisor might prefer:

$$\square (R(d1) \wedge Sub^*(d1, d2) \rightarrow \diamond_{\leq k} R(d2))$$

where  $\diamond_{\leq k}$  is “eventually within  $k$  time steps”. The value of  $k$  is operationally meaningful and controller-declared; a defensible default is  $k = 1$  revocation-propagation cycle (roughly 1 second in a typical agent runtime).

**F.3.2 BlocksNewWork as the default post-revocation state.** Once revoked, no new work unit may be accepted under that delegation.

$$\square (R(d) \rightarrow \square \neg \exists t. New(t, d))$$

The inner  $\square$  is essential: the property must hold in all subsequent states, not just the next one.

**F.3.3 CancelsQueuedWork.** Once revoked, all queued (not yet executing) work under the delegation is cancelled.

$$\square (R(d) \rightarrow \circ \forall t. \neg Q(t, d))$$

Immediate rather than eventual: queued work cancels in the next state after revocation.

**F.3.4 StopsInflightWork as opt-in.** In-flight work stops only if the opt-in applies; otherwise, in-flight work runs to completion (the default in §6.1).

$$\begin{aligned} &\square (R(d) \wedge Opt(StopsInflight, d) \rightarrow \circ \forall t. W(t, d) \rightarrow \neg InFlight(t)) \\ &\square (R(d) \wedge \neg Opt(StopsInflight, d) \rightarrow \forall t. W(t, d) \rightarrow (InFlight(t) U Completed(t))) \end{aligned}$$

The first formula: under opt-in, in-flight stops in the next state. The second formula: without opt-in, in-flight work continues executing until it completes (Until semantics).

**F.3.5 MarksArtifactsStale.** Artefacts produced under the revoked authority are eventually marked stale, and this marking is persistent.

$$\begin{aligned} &\square (R(d) \wedge ProducedUnder(x, d) \rightarrow \diamond Stale(x, d)) \\ &\square (Stale(x, d) \rightarrow \square Stale(x, d)) \end{aligned}$$

Persistence is a safety property: once stale-marked, an artefact does not revert to fresh.

## F.4 What these formulas establish and what they do not

---

Would establish, conditional on (i) the transition system in `mcs_revocation_model_v0_2.tla` being a faithful model of a concrete revocation pipeline, and (ii) successful TLC model-checking of the LTL formulas in §F.3 against bounded instantiations: that the five §6.1 effects are mutually consistent; that transitive revocation does not deadlock; that the `StopsInflightWork` opt-in is correctly disjoint from the default; that staleness is persistent. Neither (i) nor (ii) is satisfied at v0.5.1 release.

Do not establish: that any concrete revocation pipeline implements the transition system correctly; that the bounded-latency value of  $k$  is operationally achievable in production; that revocation semantics under network partitions or Byzantine failure match the LTL specification (both are beyond propositional LTL and would require either probabilistic or real-time temporal logic); that the LTL formulas capture the full revocation intent of the GDPR Art. 17 right-to-erasure use case or the NIS2 Art. 23 incident-response use case.

## F.5 Verification status

---

The TLA+ specification has not yet been model-checked against a concrete model. Model-checking with TLC requires a bounded finite-state abstraction of the delegation graph, the work queue, and the artefact set; the bound choices are non-trivial and are the first task for model-checker execution. Publication of verification transcripts is committed for a later release of the `mcs_revocation_model` artefact, independent of the paper timeline.

## F.6 Why LTL rather than CTL or real-time logics

---

LTL is sufficient for the five effects because each is expressible as a property over linear traces. CTL would be required for branching-time properties (for example, “along some path, revocation propagates within  $k$  steps; along another path, it propagates within  $k+1$ ”); the §6.1 specification does not require branching semantics. Real-time logics (TCTL, MITL) would be required for properties with hard timing constraints (for example, “revocation propagates within exactly 100ms”); the declared operational default ( $k \approx 1$  second) is loose enough that LTL suffices. Extensions to TCTL are natural if a future release needs to specify deadline-sensitive revocation for safety-critical AI subsystems.

---

End of paper.

© 2026 J. Janssen Apparens

## 2. Framework specification

---

# MCS Framework Specification

**Companion artefact** to *A Supervisory-Evidence Ontology for Agentic AI under EU Law: Minimum Conceptual Set and Temporal Extension* (Janssen, April 2026).

**Filename:** `mcs_framework_specification.md` **Purpose:** human-readable equivalent of `mcs_vocabulary.ttl`. Documents the 23 MCS classes, their properties, regulatory anchors, and reuse pattern for readers without RDF tooling. **Licence:** CC BY 4.0 **Version:** derived from `mcs_vocabulary.ttl` v0.2 (25 April 2026)

## §1. Reading guide

This specification documents the 23 classes of the Minimum Conceptual Set (MCS) defined in §4.3 of the companion paper. The authoritative artefact is the OWL/RDF Turtle vocabulary at `mcs_vocabulary.ttl`. This document provides:

- Slot-by-slot class definitions with operational semantics.
- Regulatory anchors per class.
- Property declarations with domain, range, and rationale.
- Cross-references to formal sketches F1-F16 (Appendix A of the paper).
- Cross-references to SHACL profile shapes in `mcs_profiles_v1_2.ttl` and `mcs_sme_profile_v0_1.ttl`.

The paper's §2.5 counting convention applies: 23 conceptual slots / 23 net-new MCS classes, with 5 DPV-reuse slots (6-10) and 2 PROV-O-direct-reuse slots (3, 19) producing 0 net new classes each. Subtype hierarchies under a single root are counted as one MCS slot. The machine-readable Turtle file necessarily contains additional `owl:Class` declarations for subtypes, supporting classes, and named categories; the generated manifest reports the raw OWL count separately.

## Architectural layering

The 23 classes are split into two layers per §2.5 and §4.3 of the paper:

- **mcs-core** (slots 1-12): agent-behaviour layer. Twelve classes covering agent action primitives, decision and intervention semantics, and delegation-and-policy temporal structure.
- **mcs-supervisory** (slots 13-23): supervisory-evidence layer. Eleven classes covering sectoral regime pointers, supervisory scope, typed log events, monitoring observations, drift indicators, evidence artefacts, serious incidents, and the NonQualifyingADM residual.

A planned split into `mcs_core.ttl` and `mcs_supervisory.ttl` will operationalise the layering as distinct `owl:imports` targets. The current consolidated file is `mcs_vocabulary.ttl`.

## Disclaimer

Structural-completeness vocabulary only. Validation against shapes derived from this vocabulary yields structural conformance with the named profile's structural constraints; it does not establish compliance with the AI Act, the GDPR, NIS2, DORA, the CRA, or any other instrument of Union or national law. Structural completeness is necessary but not sufficient for supervisory documentation obligations.

## §2. Namespace and prefix conventions

Prefix	IRI	Status
<code>mcs:</code>	<code>http://apparens.nl/mcs#</code>	MCS ontology root
<code>gov:</code>	<code>http://apparens.nl/mcs/gov#</code>	MCS class and property namespace
<code>prov:</code>	<code>http://www.w3.org/ns/prov#</code>	PROV-O (W3C Recommendation 2013); hard import
<code>time:</code>	<code>http://www.w3.org/2006/time#</code>	OWL-Time (W3C Recommendation 2017); hard import
<code>dpv:</code>	<code>https://w3id.org/dpv#</code>	Data Privacy Vocabulary v2.3; hard import
<code>gdprov:</code>	<code>http://openscience.adaptcentre.ie/ontologies/GDPROv#</code>	GDPROv; soft import
<code>eu-aiact:</code>	<code>https://w3id.org/dpv/legal/eu/aiact#</code>	DPV EU-AIAct extension v2.3; soft import
<code>owl:</code> , <code>rdf:</code> , <code>rdfs:</code> , <code>xsd:</code> , <code>skos:</code> , <code>dcterms:</code>	W3C standard	utility

Hard imports are operative `owl:imports` targets. Soft imports are referenced via prefix where available; their canonical IRIs are not consistently resolvable as of April 2026. The MCS uses soft-imported terms by name (e.g., `eu-aiact:Provider`) for documentation but does not block on their operational availability.

## §3. mcs-core: agent-behaviour layer (slots 1-12)

### Slot 1: gov:Agent (with subtypes SoftwareAgent, HumanAgent, OrganizationalAgent)

**Definition:** Abstract supervisory-scope agent. Covers software agents (LLM-based or symbolic), human agents acting in operational capacity, and organisational agents (controllers, processors, deployers).

**Subclass of:** `prov:Agent` (PROV-O reuse for attribution semantics).

**Subtypes:** - `gov:SoftwareAgent` : subclass of `gov:Agent` and `prov:SoftwareAgent` . Autonomous or semi-autonomous software process, including LLM-based agents. - `gov:HumanAgent` : subclass of `gov:Agent` and `prov:Person` . Natural person acting in operational capacity within a controller or processor. - `gov:OrganizationalAgent` : subclass of `gov:Agent` and `prov:Organization` . Legal entity acting as controller, processor, deployer, or other AI Act/GDPR role-bearer.

**Regulatory anchor:** GDPR Art. 4(7)–(8); AI Act Art. 3(3)–(4); SCHUFA paragraph 48.

**Reuse pattern:** PROV-O present (formal class declaration); DPV partial; agent ontologies partial-to-present per Table 4.1 in the paper.

**Cross-reference:** paper Table 4.3a slot 1.

---

### Slot 2: gov:LegalRole (with eight AI Act and GDPR subtypes)

**Definition:** Abstract legal role under GDPR or AI Act. Time-indexed attribution to `gov:Agent` via `gov:RoleAssignment` . A single organisation can simultaneously hold `dpv:Controller` under GDPR and `gov:Deployer` under the AI Act; the assignments are distinct.

**Subtypes:** - `gov:Provider` (AI Act Art. 3(3)) - `gov:Deployer` (AI Act Art. 3(4)) - `gov:Distributor` (AI Act Art. 3(7)) - `gov:Importer` (AI Act Art. 3(6)) - `gov:Manufacturer` (AI Act Art. 3(8)) - `gov:AuthorisedRepresentative` (AI Act Art. 3(5)) - `gov:Controller` (GDPR Art. 4(7); see-also `dpv:Controller`) - `gov:Processor` (GDPR Art. 4(8); see-also `dpv:Processor`)

**Supporting class:** `gov:RoleAssignment` (subclass of `prov:Attribution`) . Required properties: - `gov:hasAgent` (object property): the agent receiving the role. - `gov:hasRole` (object property): the role being assigned. - `gov:assignmentInterval` (object property, range `time:ProperInterval`): time-bounded validity.

**Regulatory anchor:** GDPR Arts 4(7)–(8), 26, 28; AI Act Art. 3(3)–(8).

**Reuse pattern:** DPV present for GDPR roles (`dpv:Controller`, `dpv:Processor`); AI Act roles absent in surveyed agent ontologies.

**Cross-reference:** paper Table 4.3a slot 2; §4.2 Legal Role Triad.

---

### Slot 3: gov:Activity (PROV-O reuse, zero net new)

**Definition:** Agent action within MCS supervisory scope.

**Subclass of:** `prov:Activity` . Inherits `prov:startedAtTime`, `prov:endedAtTime`, `prov:wasAssociatedWith` .

**Regulatory anchor:** AI Act Art. 12.

**Reuse pattern:** Direct PROV-O reuse with no new properties. Counted as 0 net new classes under §2.5 convention.

**Cross-reference:** paper Table 4.3a slot 3.

---

### Slot 4: gov:ToolInvocation

**Definition:** Specific tool-call activity performed by an agent. Anchored in AI Act Art. 13 transparency obligations.

**Subclass of:** `gov:Activity` .

**Properties:** - `gov:hasTool` (object property): reference to the tool, function, or external service invoked. - `gov:hasInvocationInput` (object property): input arguments to the tool call. - `gov:hasInvocationOutput` (object property): output of the tool call.

**Regulatory anchor:** AI Act Art. 13; GDPR Art. 22 by tool-chain extension (paper §5.3 sketch F3); NIS2 Art. 23(3) for tool invocations contributing to significant incidents (paper §5.3 sketch F4).

**Reuse pattern:** Agent Spec partial; surveyed vendor ontologies partial-to-present.

**Cross-reference:** paper Table 4.3a slot 4; mapping matrix cells C01–C06.

---

### Slot 5: gov:MemoryEvent (with subtypes MemoryRead, MemoryWrite)

**Definition:** Provenance of model or agent memory access; abstract over read and write.

**Subclass of:** `gov:Activity` .

**Subtypes:** - `gov:MemoryRead` : an agent reads from memory. Source for paper §5.3 sketch F5 (research conjecture on Art. 22 contributing factor). - `gov:MemoryWrite` : an agent writes to memory. Subject to NIS2 Art. 23(4) preservation requirement during incident windows; see App A.15 sketch F12.

**Properties:** - `gov:hasMemoryReference` (datatype property, range `xsd:anyURI`): URI or addressable identifier for the memory location.

**Regulatory anchor:** AI Act Art. 12 (logging); AI Act Art. 72 (post-market monitoring); NIS2 Art. 23(4) (incident-window preservation).

**Reuse pattern:** Atlan and Skan present; other agent ontologies partial.

**Cross-reference:** paper Table 4.3a slot 5; mapping matrix cells C07–C18.

---

#### Slot 6: `dpv:DataSubject` (DPV reuse, zero net new)

**Definition:** Data subject reference per GDPR.

**Reuse:** imported directly via `owl:imports` from DPV v2.3. Not redeclared in MCS namespace.

**Regulatory anchor:** GDPR Art. 4(1), 15, 22.

**Cross-reference:** paper Table 4.3a slot 6.

---

#### Slot 7: `dpv:PersonalData` (DPV reuse, zero net new)

**Definition:** Personal-data category hierarchy including Art. 9 special categories.

**Reuse:** imported directly from DPV v2.3.

**Regulatory anchor:** GDPR Art. 9, 10, 30(1)(c).

**Cross-reference:** paper Table 4.3a slot 7.

---

#### Slot 8: `dpv:Purpose` (DPV reuse, zero net new)

**Definition:** Purpose of processing per GDPR.

**Reuse:** imported directly from DPV v2.3.

**Regulatory anchor:** GDPR Art. 5(1)(b), 6(4), 30.

**Cross-reference:** paper Table 4.3a slot 8.

---

#### Slot 9: `dpv:LegalBasis` (DPV reuse, zero net new)

**Definition:** Lawful basis for processing per GDPR.

**Reuse:** imported directly from DPV v2.3.

**Regulatory anchor:** GDPR Art. 6, 9.

**Cross-reference:** paper Table 4.3a slot 9.

---

#### Slot 10: `dpv:Consent plus GDPRov consent lifecycle` (DPV plus GDPRov reuse)

**Definition:** Consent record with full lifecycle (given, withdrawn, expired, renewed).

**Reuse:** `dpv:Consent` from DPV v2.3 plus GDPRov consent-lifecycle vocabulary.

**Regulatory anchor:** GDPR Art. 6(1)(a), 7.

**Cross-reference:** paper Table 4.3a slot 10.

---

#### Slot 11: `gov:Decision` (with three subtypes plus `gov:Score plus drawsStronglyOn`)

**Definition:** Decision concerning a data subject. Three-way taxonomy anchored in GDPR Art. 22(1) and SCHUFA upstream-attribution path.

**Subtypes:** - `gov:SolelyAutomatedDecision` : decision based solely on automated processing within Art. 22(1) scope. Profile-dependent boundary. - `gov:AssistedDecisionWithMeaningfulHumanIntervention` : decision with human intervention meeting WP251rev.01 plus AP handvatten conjunctive criteria. - `gov:NonAutomatedDecision` : decision without ADM substrate; outside Art. 22 scope by absence of automated processing.

**Supporting class:** `gov:Score`. A preparatory automated probability or risk score. Per SCHUFA para 73, a score is itself a decision under Art. 22(1) where downstream actors draw strongly on it.

**Predicate:** `gov:drawsStronglyOn` (object property, domain `gov:Decision`, range `gov:Score`). Profile-dependent operationalisation per paper §6.4: - Profile A (AP-permissive, quantitative): flip-rate at  $\varepsilon=0.10$  exceeds  $M=0.50$ . - Profile B (CNIL-strict, qualitative): any of three judicial indicia (specialised scoring model; refusal frequency  $\geq 0.90$ ; absence of documented intervention).

**Properties on `gov:Decision`:** - `gov:hasLegalEffect` (datatype property, range `xsd:boolean`): Art. 22(1) qualifier; produces legal effects concerning the data subject or similarly significantly affects them. - `gov:hasDocumentedIntervention` (datatype property, range `xsd:boolean`): Profile B qualitative indicium 3. - `gov:refusalFrequencyGivenScore` (datatype property, range `xsd:decimal`): Profile B qualitative indicium 2; operationalises SCHUFA para 48 “in almost all cases”.

**Properties on `gov:Score`:** - `gov:flipRateAtEpsilon10` (datatype property, range `xsd:decimal`): Profile A quantitative property. - `gov:usesSpecialisedScoringModel` (datatype property, range `xsd:boolean`): Profile B qualitative indicium 1; Wiesbaden VG January 2026.

**Regulatory anchor:** GDPR Art. 22(1); SCHUFA paragraph 73; Dun and Bradstreet paragraph 40.

**Reuse pattern:** absent in all surveyed ontologies.

**Cross-reference:** paper Table 4.3a slot 11; §5.3 sketch F1; §6.4 dual operationalisation; §4.6 Findings 3 and 4.

---

## Slot 12: `gov:HumanIntervention`

**Definition:** Six-property schema operationalising the WP251rev.01 plus AP handvatten meaningful-intervention test, with the `authorityActuallyExercised` property added per Uber/Ola Hof Amsterdam (case B10, ECLI:NL:GHAMS:2023:793). All six properties must hold for an intervention to count as meaningful under the conjunctive test.

**Properties:** - `gov:hasAuthorityToDeviate` (datatype property, range `xsd:boolean`): reviewer holds documented authority to override the automated outcome. - `gov:hasUnderstandingOfLogic` (datatype property, range `xsd:boolean`): reviewer holds documented understanding of the system’s decision logic per Art. 15(1)(h). - `gov:consideredAllRelevantData` (datatype property, range `xsd:boolean`): reviewer documented consideration of relevant case-specific factors. - `gov:hasMarginOfDiscretion` (datatype property, range `xsd:boolean`): reviewer’s role permits substantive discretion, not merely rubber-stamping. - `gov:interventionDurationSeconds` (datatype property, range `xsd:nonNegativeInteger`): time spent on intervention. Default operational threshold  $\geq 60$ s for non-symbolic review, sector-calibrated (5min for complex clinical triage; 15s for high-volume fraud-detection alerts conditional on case-specific examination). See paper §2.7. - `gov:authorityActuallyExercised` (datatype property, range `xsd:boolean`): per Uber/Ola Hof Amsterdam (B10), authority to deviate is necessary but not sufficient. The authority must be actually exercised in practice; population-level deviation rate  $> \theta$  (illustrative  $\theta=10\%$ ). Profile B requires this property=true for MeaningfulIntervention.

**Relation property:** - `gov:contributesTo` (object property, domain `gov:HumanIntervention`, range `gov:Decision`).

**Regulatory anchor:** GDPR Art. 22(3); AI Act Art. 14; WP251rev.01; AP handvatten July 2025.

**Reuse pattern:** absent in all surveyed ontologies.

**Cross-reference:** paper Table 4.3a slot 12; §4.6 Finding 3 (Uber/Ola jurisprudence); App A.12 sketch F9.

---

## §4. mcs-supervisory: supervisory-evidence layer (slots 13-23)

---

### Slot 13: `gov:Policy` and `gov:PolicyVersion`

**Definition:** An organisational or technical policy governing agent behaviour, plus its versioned snapshots.

**Classes:** - `gov:Policy`: the policy as identity (versionless reference). - `gov:PolicyVersion`: subclass of `prov:Entity`. Versioned snapshot with temporal validity and supersession.

**Properties on `gov:PolicyVersion`:** - `gov:supersedesPolicy` (object property, range `gov:PolicyVersion`). - `gov:isMorePermissive` (datatype property, range `xsd:boolean`): used in App A.10 sketch F7 (lex mitior conjecture, RC-flagged). - `gov:policyValidityInterval` (object property, range `time:ProperInterval`).

**Regulatory anchor:** AI Act Art. 9 risk management; Art. 17 QMS.

**Cross-reference:** paper Table 4.3b slot 13; mapping matrix cells C31-C36; PolicyVersionShape in `mcs_profiles_v1_2.ttl` (non-overlapping validity intervals across successor versions).

---

### Slot 14: `gov:RiskClassification` (with four AI Act subtypes)

**Definition:** AI Act risk classification of an AI system.

**Subtypes:** - `gov:ProhibitedAI`: AI Act Art. 5 prohibited practice. - `gov:HighRiskAI`: AI Act Art. 6 plus Annex III high-risk system. Subject to Chapter III obligations. - `gov:LimitedRiskAI`: AI Act Art. 50 transparency obligations only. - `gov:MinimalRiskAI`: outside the scope of Arts. 5, 6, and 50; voluntary

codes of conduct under Art. 95.

**Regulatory anchor:** AI Act Art. 5, 6, 50, Annex III; ISO/IEC 23894.

**Reuse pattern:** ISO/IEC partial (prose); absent as formal class elsewhere.

**Cross-reference:** paper Table 4.3b slot 14; App A.13 sketch F10.

---

### Slot 15: gov:IntendedPurpose

**Definition:** AI Act Art. 3(12): use for which the AI system is intended by the provider, including specific context and conditions of use, as specified in instructions for use. Distinct from `dpv:Purpose` (which is the GDPR Art. 5(1)(b) purpose-of-processing concept).

**Regulatory anchor:** AI Act Art. 3(12), 13.

**Reuse pattern:** JTC 21 partial (prose); Agent Spec Intent partial.

**Cross-reference:** paper Table 4.3b slot 15.

---

### Slot 16: gov:SectoralRegime (with named instances)

**Definition:** Pointer to a sectoral regulatory regime applicable to the AI system.

**Properties:** - `gov:competentSupervisor` (object property, range `gov:SupervisoryScope`).

**Named instances:** - `gov:MiFID_II`: Directive 2014/65/EU on markets in financial instruments. - `gov:MDR`: Regulation (EU) 2017/745 on medical devices. Rule 11 governs autonomous clinical-decision agents. - `gov:CRR_CRD`: Capital Requirements Regulation 575/2013 plus Capital Requirements Directive 2013/36/EU. - `gov:IDD`: Insurance Distribution Directive 2016/97. - `gov:AWB`: Algemene wet bestuursrecht (Dutch General Administrative Law Act).

**Regulatory anchor:** sectoral statutes.

**Reuse pattern:** absent in all surveyed ontologies.

**Cross-reference:** paper Table 4.3b slot 16; §7.6 regulatory-coherence assumption.

---

### Slot 17: gov:SupervisoryScope (ten Dutch supervisors plus EU bodies)

**Definition:** Competent authority with defined supervisory mandate over an AI system or processing activity.

**Named instances (Dutch supervisors, ten roles, nine entities):** - `gov:AFM`: Autoriteit Financiële Markten. - `gov:DNB`: De Nederlandsche Bank. - `gov:IGJ`: Inspectie Gezondheidszorg en Jeugd. - `gov:AP`: Autoriteit Persoonsgegevens; coordinating algorithm supervisor since February 2025. - `gov:RDI`: Autoriteit Digitale Infrastructuur. - `gov:NLA_SZW`: Nederlandse Arbeidsinspectie. - `gov:NCSC`: Nationaal Cyber Security Centrum. - `gov:ILT`: Inspectie Leefomgeving en Transport. - `gov:ACM_CM`: Autoriteit Consument en Markt, consumer-and-markets mandate. - `gov:ACM_Energy`: ACM energy division.

**Named instances (EU-level bodies):** - `gov:EDPB`: European Data Protection Board. - `gov:AIOffice`: European Commission AI Office.

**Regulatory anchor:** AI Act Art. 70; AP-RDI Final advice 7 November 2024.

**Cross-reference:** paper Table 4.3b slot 17; §5.2 ten-supervisor model.

---

### Slot 18: gov:DelegationGrant (with RevocationEvent, ValidityInterval, AuthorityScope)

**Definition:** Temporal class with validity windows and revocation semantics. Subclass of `prov:Delegation` extending it with explicit validity intervals and authority scope.

**Classes:** - `gov:DelegationGrant`: subclass of `prov:Delegation`. Required properties `gov:hasValidityInterval`, `gov:hasAuthorityScope`. - `gov:RevocationEvent`: subclass of `prov:Activity`. Event revoking a `DelegationGrant`. Triggers up to five effects per §6.1. - `gov:ValidityInterval`: subclass of `time:ProperInterval`. Time-bounded validity for a delegation, role assignment, or policy version. - `gov:AuthorityScope`: the set of actions, resources, or data the delegated authority covers.

**Properties:** - `gov:hasValidityInterval` (object property, range `gov:ValidityInterval`). - `gov:hasAuthorityScope` (object property, range `gov:AuthorityScope`). - `gov:isSubdelegationOf` (object property, domain and range `gov:DelegationGrant`). Subdelegation chain. App A.9 sketch F6 specifies transitive revocation per GDPR Art. 28(2),(4).

**Five revocation effects (§6.1):** - `gov:BlocksNewWork`: default effect; no new work accepted under revoked delegation. - `gov:CancelsQueuedWork`: queued (not yet executing) work under the delegation is cancelled. - `gov:StopsInflightWork`: opt-in effect; in-progress tasks are stopped. Operationally risky default; specified as opt-in. - `gov:RequiresReview`: tasks completed recently under revoked authority must be reviewed before acceptance downstream. - `gov:MarksArtifactsStale`: default effect; artefacts produced under revoked authority marked stale; persistence is a safety property (App F.3.5).

**Properties on gov:Decision related to revocation:** - gov:executedUnderRevokedAuthority (datatype property, range xsd:boolean): default false. True requires explicit policy authorisation. - gov:cachedUnderRevokedAuthority (datatype property, range xsd:boolean): artefact produced before revocation but accessed after; lawfully produced but may need review.

**Regulatory anchor:** GDPR Art. 28; AI Act Art. 25.

**Reuse pattern:** ACP Framework present (cryptographic delegation); PROV-O partial (basic delegation only).

**Cross-reference:** paper Table 4.3b slot 18; §6.1 revocation propagation; App F LTL formalisation; mapping matrix cells C19–C30.

---

### Slot 19: gov:ProvenanceChain (PROV-O reuse, zero net new)

**Definition:** Integration point reusing prov:wasDerivedFrom, prov:wasInformedBy, prov:wasInfluencedBy.

**Reuse:** counted as 0 net new classes under §2.5 counting convention (direct PROV-O reuse).

**Regulatory anchor:** AI Act Art. 10; Art. 12.

**Cross-reference:** paper Table 4.3b slot 19.

---

### Slot 20: gov:LogEvent (typed AI Act subclasses)

**Definition:** Typed log events at statutory-subclass granularity. AI Act Art. 12(2) lists five categories implicitly; the MCS makes them first-class subclasses.

**Subclass of:** prov:Activity.

**Subtypes:** - gov:RiskIdentificationLogEvent : AI Act Art. 12(2)(a); situations that may result in the AI system presenting a risk. - gov:PostMarketMonitoringLogEvent : AI Act Art. 12(2)(b); post-market monitoring substantial modification events. - gov:OperationMonitoringLogEvent : AI Act Art. 12(2)(c); functional operation monitoring. - gov:BiometricIdentificationLogEvent : AI Act Art. 12(3); mandatory fields for biometric identification systems. - gov:IncidentLogEvent : AI Act Art. 73 plus NIS2 Art. 23 incident reporting events.

**Regulatory anchor:** AI Act Art. 12(2)–(3); Art. 26(6); NIS2 Art. 23.

**Reuse pattern:** generic logging ubiquitous; typed-to-statute subclasses absent in surveyed ontologies.

**Cross-reference:** paper Table 4.3b slot 20.

---

### Slot 21: gov:MonitoringObservation

**Definition:** Subclass of prov:Entity recording a metric observation in post-market monitoring context.

**Properties:** - gov:observedMetric (datatype property). - gov:observedAtTime (datatype property, range xsd:dateTime). - gov:monitoringPurpose (object property).

**Regulatory anchor:** AI Act Art. 72; ECB Guide to Internal Models; SAFEST framework (DNB 2019).

**Reuse pattern:** absent in all surveyed ontologies.

**Cross-reference:** paper Table 4.3b slot 21.

---

### Slot 22: gov:DriftIndicator (with four subtypes)

**Definition:** Indicator of drift in data, concept, model performance, or control.

**Subtypes:** - gov:DataDrift : distributional shift in input data. - gov:ConceptDrift : shift in target-variable distribution or relationship between inputs and outputs. - gov:ModelPerformanceDrift : degradation in model performance metrics over time. - gov:ControlDrift : erosion of operational controls (oversight cadence, intervention rates) over time.

**Properties:** - gov:driftMagnitude (datatype property, range xsd:decimal). - gov:driftDetectedAt (datatype property, range xsd:dateTime). - gov:triggersRefresh (datatype property, range xsd:boolean): whether the drift indicator triggers evidence-refresh per §6.2 Bayesian decay framework.

**Regulatory anchor:** AI Act Art. 72 continuous performance analysis; PCAOB AS 1105.

**Reuse pattern:** absent in all surveyed ontologies.

**Cross-reference:** paper Table 4.3b slot 22.

---

### Slot 23: gov:SupervisoryEvidenceArtifact + gov:Incident + gov:NonQualifyingADM (one slot, three thematic classes)

**Definition (a) gov:SupervisoryEvidenceArtifact:** Evidence artefact with source hierarchy and probative weight.

**Subclass of:** `prov:Entity`.

**Source hierarchy** (PCAOB AS 1105 illustrative defaults, controller-configurable; see App A.11 sketch F8): - `gov:SystemLogSource` : default reliability weight 0.85. - `gov:HumanAttestationSource` : default reliability weight 0.60. - `gov:ExternalReportSource` : default reliability weight 0.70. - `gov:ThirdPartyAuditSource` : default reliability weight 0.90.

**Properties:** - `gov:probativeWeight` (datatype property, range `xsd:decimal` in [0.0, 1.0]). Bayesian decay function with sectoral half-life priors per §6.2: - ISAE 3402: 12 months. - ECB CRR/CRD: 12 months. - MDR: 6 months. - AI Act Art. 12: 6 months (floor, not ceiling). - AI Act Art. 72: 12 months. - NIS2 Art. 21: 6 months. - Manual attestation: 3 months. - `gov:initialWeight` (datatype property, range `xsd:decimal`). - `gov:lastRefreshedAt` (datatype property, range `xsd:dateTime`).

**Definition (b) gov:Incident:** Serious incident under AI Act Art. 73 or NIS2 Art. 23. Significance assessed per F4 sketch (paper §5.3).

**Subclass of:** `prov:Activity`.

**Properties:** - `gov:detectionTime` (datatype property, range `xsd:dateTime`). - `gov:incidentWindow` (object property, range `time:ProperInterval`): temporal window within which delegation events and memory writes must be preserved per NIS2 Art. 23(4).

**Definition (c) gov:NonQualifyingADM:** Residual class for ADM enforcement decided on GDPR provisions other than Art. 22.

**Empirical anchor:** paper §4.6 Finding 2; six of 25 cases in `mcs_case_sample_v0_2.csv`. Sensitivity range [4%, 52%] under borderline-case flipping; Wilson 95% CI [11.5%, 43.4%]. The category exists and is structurally significant; the proportion is empirically contingent.

**Subtypes:** - `gov:Art5FairnessADM` : ADM enforcement decided on GDPR Art. 5(1)(a) fairness ground. Empirical anchor: case B20 Belastingdienst Toeslagen. - `gov:Art6LawfulBasisADM` : ADM enforcement decided on GDPR Art. 6 lawful-basis ground. Empirical anchor: cases B18 Clearview AI multi-jurisdictional, B24 CNIL Clearview fine. - `gov:Art9SpecialCategoryADM` : ADM enforcement decided on GDPR Art. 9 special-category-data ground. Empirical anchor: case B19 live FR in schools. - `gov:Art25ByDesignADM` : ADM enforcement decided on GDPR Art. 25 by-design ground. Empirical anchor: case B17 Norwegian IB grading. - `gov:Art35DPIAADM` : ADM enforcement decided on GDPR Art. 35 DPIA-failure ground. Plus case B22 SyRI decided on Art. 8 ECHR proportionality (cognate residual).

**Regulatory anchor:** AI Act Art. 26(10); AI Act Art. 73; NIS2 Art. 23; GDPR Arts 5(1)(a), 6, 9, 25, 35; Art. 8 ECHR.

**Reuse pattern:** absent in all surveyed ontologies.

**Cross-reference:** paper Table 4.3b slot 23; §4.6 Findings 2; §6.2 evidence decay; §6.3 SHACL shapes; mapping matrix cells C18, C24, C30 (incident windows).

---

## §5. SHACL profile shapes (companion artefact `mcs_profiles_v1_2.ttl`)

---

The vocabulary supports SHACL validation through the companion shapes file `mcs_profiles_v1_2.ttl`. The file declares:

- `gov:SupervisoryProfile` abstract class with two operational profiles:
  - **Profile A** (AP-permissive, quantitative): `drawsStronglyOn` evaluated via flip-rate at  $\epsilon=0.10$ ,  $M=0.50$ .
  - **Profile B** (CNIL-strict, qualitative): `drawsStronglyOn` evaluated via three judicial indicia.
- Profile C (UK ICO post-Brexit) and Profile D (Italian Garante post-Law 132/2025) ship as stubs.
- Fourteen core temporal-extension shapes (§6.3) plus seven profile-layer shapes (§5.5): `ProfileA_Art22ScopeShape`, `ProfileA_DrawsStronglyOnShape`, `ProfileB_Art22ScopeShape`, `ProfileB_DrawsStronglyOnShape`, `ProfileB_MeaningfulInterventionShape`, `CrossProfileGovernanceShape`, `ProfileDeltaShape`. Twenty-one `NodeShapes` total in `mcs_profiles_v1_2.ttl`. The fourteen core shapes are: `DelegationGrantShape`, `PolicyVersionShape`, `EvidenceDecayShape`, `EvidenceQualityShape`, `MonitoringObservationShape`, `DriftIndicatorShape`, `AgentActionShape`, `RoleAssignmentShape`, `BiometricLogShape`, `DecisionShape`, `SchufaUpstreamShape`, `HumanInterventionShape` (added v0.5.1), `RevocationEventShape` (added v0.5.1), and `LogEventSubclassGranularityShape` (added v0.5.1). The three v0.5.1-added shapes resolve previously dangling references in `mcs_sme_profile_v0_2.ttl`.
- `gov:CrossProfileGovernanceShape` and `gov:ProfileDeltaShape` for surfacing profile divergences as governance artefacts.

For SME-eligible controllers, `mcs_sme_profile_v0_1.ttl` provides a parallel proportionality profile with structured prose substitution for a narrow set of shapes (§5.6).

---

## §6. Cross-references to companion artefacts

Artefact	Function
<code>mcs_vocabulary.ttl</code>	Authoritative OWL/RDF vocabulary (this specification)

	documents).
mcs_profiles_v1_2.ttl	SHACL profile shapes (Profiles A and B; C and D stubs).
mcs_sme_profile_v0_1.ttl	SME proportionality profile per §5.6.
mcs_revocation_model_v0_1.tla	TLA+ sketch for revocation propagation, pending TLC verification (App F).
report_translator_v0_1.py	Plain-language ValidationReport translator (App E.7).
mcs_case_sample_v0_2.csv	25-case empirical sample with Step 1-5 protocol coding.
mcs_mapping_matrix_v0_2.csv	48-cell mapping matrix (8 primitives × 6 legal categories).
coding_rubrics_v0_2.md	Reproducibility protocols for Track 1.
validation_pack_v1.md	Pre-registered open-call mechanics for three validation tracks.
mcs_scenarios_v0_2.md	Seven worked scenarios stress-testing the MCS internally.

## §7. Counting convention summary (paper §2.5)

Counting category	Slots	Net-new classes
mcs-core (agent-behaviour)	1, 2, 4, 5, 11, 12 (6 slots with new classes)	6
mcs-core (PROV-O reuse)	3	0
mcs-core (DPV reuse)	6, 7, 8, 9, 10	0
<b>mcs-core total</b>	<b>slots 1 to 12</b>	<b>6 (net-new in core)</b>
mcs-supervisory (new classes)	13, 14, 15, 16, 17, 18, 20, 21, 22, 23	10
mcs-supervisory (PROV-O reuse)	19	0
<b>mcs-supervisory total</b>	<b>slots 13 to 23</b>	<b>10 (net-new in supervisory)</b>
<b>MCS total (slot count)</b>	<b>23 slots</b>	n/a
<b>MCS total (net-new classes, strict reuse-zero accounting)</b>	n/a	<b>16 net-new</b>

The figure 23 is a slot count. The figure 16 is a net-new-class count under strict reuse-zero accounting (5 DPV reuses in slots 6 to 10 plus 2 PROV-O reuses in slots 3 and 19, each contributing 0 net-new classes). The figure 69 is the raw owl:Class count in the Turtle source `mcs_vocabulary.ttl` once subtypes, support classes, and named instances are counted. Under alternative conventions (each thematic class in slot 23 counted separately, subtype hierarchies flattened) the count ranges from approximately 30 to 43 classes. Conflating the slot count with the class count produces the inconsistency previously noted in earlier drafts.

## §8. Versioning and licensing

- **Vocabulary file:** `mcs_vocabulary.ttl`, current version derived from paper §4.3 (April 2026).
- **Specification file:** `mcs_framework_specification.md` (this document).
- **Licence:** CC BY 4.0 for both the vocabulary and the specification.
- **Provenance:** registered as part of the Apparens MCS specification suite. The MCS is a separate work from the AI Enterprise Control Index (i-DEPOT 158508 at BOIP) and is not covered by that registration. Anchor for the MCS specification suite is Zenodo DOI [10.5281/zenodo.19758441](https://doi.org/10.5281/zenodo.19758441) (v0.5.1).

End of `mcs_framework_specification.md`. © 2026 J. Janssen, Apparens. CC BY 4.0.

## Appendix A: README v0.5.1

---

Source file: `README_v0_5_1.md`

# MCS v0.5.1 release

---

This package is the v0.5.1 release of the Minimum Conceptual Set (MCS) candidate ontology for agentic AI accountability evidence under EU law. It supersedes v0.5.

**DOI (v0.5.1):** [10.5281/zenodo.19758441](https://doi.org/10.5281/zenodo.19758441) **License:** CC BY 4.0 **Cite as:** Janssen, J. (2026). A *Supervisory-Evidence Ontology for Agentic AI under EU Law: Candidate Minimum Conceptual Set and Temporal Extension* (v0.5.1) [Working paper]. Zenodo. <https://doi.org/10.5281/zenodo.19758441>

## What v0.5.1 changes from v0.5

---

v0.5.1 implements the BLOCKER and HIGH defects identified in the v0.5 adversarial review, plus the housekeeping items for D009v2, D011v2, RPS, TLC1, and DUP. See [ERRATA\\_v0\\_5\\_1.md](#) for the full reconciliation; the headlines are:

- **BLOCKER 1 (D004v2):** paper now uses three explicit figures (23 conceptual slots, 16 net-new classes under strict reuse-zero accounting, 69 owl:Class declarations) consistently across abstract, §2.5, §4.3, §6.3, §7.7, §7.8, and §9.
- **BLOCKER 2 (D003v2):** three NodeShapes added to `mcs_profiles_v1_2.ttl` (HumanInterventionShape, RevocationEventShape, LogEventSubclassGranularityShape). Total NodeShape count rises from 18 to 21. SME profile references all resolve.
- **BLOCKER 3 (S03):** scenarios renumbered into a single coherent §1 to §12 sequence with no duplicate section numbers.
- **HIGH 4 (T1B):** Track 1 blinded sample genuinely blinded; coder cannot reverse-engineer the author's classification.
- **HIGH 5 (TLA02):** TLA+ properties rewritten to quantify over the transitive closure.
- **HIGH 6 (T2P):** Track 2 runner uses `__file__`-relative paths.
- **HIGH 7 (D005v2):** §4.6 sensitivity range corrected to [16%, 36%] / [16%, 44%] anchored on  $k=6, n=25$ .
- **HIGH 8 (D008v2):** F.4 framing rewritten as conditional "Would establish"; version-string drift across paper resolved.

## Authority hierarchy

---

1. Machine-readable files are authoritative for the ontology, SHACL, CSV coding, and scripts.
2. Markdown files are human-readable specifications, scenarios, and review documentation.
3. The PDF and DOCX renderings are publication outputs derived from the markdown source. Where the rendering and the markdown disagree, the markdown wins.

## File inventory

---

### Top-level papers and specifications

- `mcs_paper_v0_5_1.md`, `.pdf`, `.docx`: candidate Minimum Conceptual Set paper.
- `mcs_framework_specification_v0_5_1.md`, `.pdf`: human-readable equivalent of the vocabulary.
- `mcs_companion_pack_v0_5_1.pdf`: single bundle PDF including all source-file appendices.

### Authoritative machine-readable source files

- `mcs_vocabulary.ttl`: 517 triples, 69 owl:Class declarations, 16 owl:ObjectProperty, 24 owl:DatatypeProperty.
- `mcs_profiles_v1_2.ttl`: 21 NodeShapes (14 core temporal-extension shapes plus 7 profile-layer shapes).
- `mcs_sme_profile_v0_2.ttl`: 3 NodeShapes (proportionality framing; legacy Exemption URIs retained).
- `mcs_revocation_model_v0_2.tla` (with v0.5.1 property rewrites) plus `.cfg`.
- `mcs_case_sample_v0_3.csv`: 25 rows, 24 columns.
- `mcs_mapping_matrix_v0_3.csv`: 48 rows, 14 columns.
- `mcs_deep_coding_matrix_v0_5.csv`: 125 rows (provisional reconstruction).
- `mcs_scenarios_v0_5_1.md`: 12 sections (renumbered).
- `coding_rubrics_v0_2.md`, `validation_pack_v1.md`.

### Reproducibility scaffolding

- `report_translator_v0_2.py`: plain-language ValidationReport translator.
- `reproduce_package_counts_v0_5_1.py`: column-name fix applied; runs cleanly.
- `shacl_examples/`: minimal pass and fail TTL examples plus expected reports plus minimal subset validator.
- `track1_blinded_package/`: redacted blinded sample plus kappa-plus-PABAK script.
- `track2_synthetic_benchmark/`: workload generator plus minimal harness with `__file__`-relative path resolution.
- `tla_model_checking/`: TLA+ v0.2 model with v0.5.1 property rewrites plus `cfg` plus `run` script with download note plus `TLC_STATUS_v0_5.md`.
- `review_notes/`: legal mapping review protocol plus 48-row review template.

### Documentation and provenance

- `README_v0_5_1.md` (this file).

- `ERRATA_v0_5_1.md` : honest reconciliation against v0.5 review and v0.4 ERRATA backlog.
- `mcs_artifact_manifest_v0_5_1.md` : with verified counts.
- `MCS_v0_5_adversarial_review.pdf` : preserved as audit-trail input.

## Validation status

---

The package is internally consistent. Validation evidence remains pre-registered, not completed:

- Track 1 inter-rater coder result: not produced. Blinded sample plus script are included.
- Track 2 publication-grade benchmark: not produced. Minimal smoke-test harness plus README disclaimer are included.
- TLC model-checking transcript: not produced. Model plus cfg plus status note are included.
- External legal mapping review: not produced. Protocol plus template are included.

## Citation

---

This release is published on Zenodo under DOI [10.5281/zenodo.19758441](https://doi.org/10.5281/zenodo.19758441). Cite the version DOI for reproducibility; cite the Concept DOI (the parent record that resolves to the latest version) when referring to the work generally.

## Versioning

---

v0.5.1 is the inaugural Zenodo release. Subsequent versions are deposited as new versions of the same record (Zenodo's "New version" workflow), each receiving its own version DOI, with the Concept DOI continuing to resolve to the latest. Planned next versions:

- v0.6: Track 1 inter-rater coder kappa-plus-PABAK results, when available.
- v0.7: TLC model-checking transcript for `mcs_revocation_model_v0_2.tla`, when available.
- v0.8: External legal mapping review of all SD/IP/RC matrix classifications, when available.

The metadata template used for v0.5.1 is preserved as `zenodo_metadata.json` in this package.

## Licence

---

CC BY 4.0 for both the vocabulary and the paper.

## Appendix B: ERRATA v0.5.1

---

Source file: `ERRATA_v0_5_1.md`

# MCS v0.5.1 Errata and Correction Backlog Closure

---

DOI (v0.5.1): [10.5281/zenodo.19758441](https://doi.org/10.5281/zenodo.19758441)

This v0.5.1 release implements the BLOCKER and HIGH defects identified in the v0.5 adversarial review. It also reconciles the v0.4 ERRATA, which overstated several fixes that had not in fact been applied to the paper or the SHACL profiles.

The intent of this document is to be honest about what is fixed, what is partially fixed, and what is deliberately deferred. Where the v0.4 ERRATA claimed a fix that was not in the source, this document marks the v0.5.1 status of the underlying issue, not the v0.4 ERRATA's claim.

---

## v0.5 review defect register reconciliation

---

The v0.5 adversarial review ( `MCS_v0_5_adversarial_review.md` ) identified BLOCKER, HIGH, and lower-severity defects against the v0.5 source pack. v0.5.1 status per defect:

### BLOCKERS

#### D004v2 (counting reconciliation: 23 net-new claim conflicts with framework spec arithmetic).

Fixed in v0.5.1. The package now uses three explicit figures everywhere: 23 conceptual slots (used wherever the paper names MCS coverage or scope); 16 net-new classes under strict reuse-zero accounting (used wherever the paper names ontology size or engineering effort); 69 owl:Class declarations (used only when discussing the source Turtle file specifically). The arithmetic: 23 slots minus 5 DPV reuses (slots 6 to 10) minus 2 PROV-O reuses (slots 3 and 19) equals 16 net-new classes (6 in mcs-core: slots 1, 2, 4, 5, 11, 12; plus 10 in mcs-supervisory: slots 13 to 18, 20 to 23).

NonQualifyingADM is the 16th net-new class. Affected passages: paper abstract, §2.5 counting convention, §4.3 opening, §6.3, §7.7, §7.8, §9 conclusion. Framework specification §7 table now has two final rows (slot count vs net-new-class count under strict reuse-zero accounting). Verification: `grep "23 net new" mcs_paper_v0_5_1.md` returns zero hits.

#### D003v2 (broken SHACL references in SME profile). Fixed in v0.5.1. Three NodeShapes added to

`mcs_profiles_v1_2.ttl`: `gov:HumanInterventionShape` (six conjunctive properties of slot 12), `gov:RevocationEventShape` (`prov:wasInformedBy` plus `prov:startedAtTime` per §6.1), and `gov:LogEventSubclassGranularityShape` (typed at AI Act statutory-subclass granularity per slot 20). Total NodeShape count in the profiles file rises from 18 to 21 (14 core temporal-extension shapes including the three v0.5.1 additions plus 7 profile-layer shapes). The SME profile's external shape references all resolve. Verification: programmatic check of regex `gov:[A-Z]\\w*Shape` in `mcs_sme_profile_v0_2.ttl` against shape declarations in `mcs_profiles_v1_2.ttl` returns an empty MISSING set.

#### S03 (scenarios numbering broken: §9 and §10 collide between original sections and new

Scenario 8/9). Fixed in v0.5.1. `mcs_scenarios_v0_5_1.md` (renamed from `_v0_3.md`) reorders so that Scenarios 8 and 9 occupy §9 and §10, and the original Cross-scenario observations and Limits sections are renumbered to §11 and §12. Total: §1 Scope, §2 to §8 Scenarios 1 to 7, §9 Scenario 8 (Hostile vendor), §10 Scenario 9 (Public-sector foundation-model), §11 Cross-scenario observations, §12 Limits. Verification: `grep "^## §" mcs_scenarios_v0_5_1.md` produces twelve unique section numbers, no duplicates.

### HIGHS

**T1B (Track 1 partial blinding leak).** Fixed in v0.5.1. `track1_blinded_case_sample_seed42_v0_5_1.csv` retains only the 10 columns a coder needs to produce an independent classification: `blind_order`, `case_id`, `jurisdiction`, `case_name`, `date`, `brief_facts`, `step1_operative_provision`, `source_citation`, `ecli_or_url`, `coder_notes`. All step-outcome columns (step1 to step5), all step-flag columns, `primary_nonqualifying_flag`, and `conditional_nonqualifying_note` are redacted. The coder must produce their own classification from the brief facts plus source material, without leakage from the author's intermediate work. The README in `track1_blinded_package/` itemises the redaction list explicitly. The kappa script (`compute_track1_kappa_v0_5_1.py`) now reports both Cohen's kappa and PABAK (Prevalence-Adjusted Bias-Adjusted Kappa) per the D032/D046 finding. Acceptance threshold: `cohen >= 0.70 AND papak >= 0.70`. Verification: programmatic check of column list returns 10 columns and 25 rows; all forbidden columns are absent.

**TLA02 (TLA+ properties under-specify transitivity).** Fixed in v0.5.1. The `BeginWork` action already used `EffectiveRevokedBy` (transitive) but the four properties  `CancelsQueuedWork`,  `StopsInflightUnderOptIn`,  `ArtefactsEventuallyStale`, and the implicit transitive form of the four-property cluster, quantified only over `revoked` directly. v0.5.1 rewrites three of those four to quantify over the transitive closure via `\\E a \\in revoked : SubStar(a, d)`. `StalenessPersists` is left direct on artefacts because the property is over the artefact itself, not its producing delegation; transitive closure is not relevant in that body. The model file remains `mcs_revocation_model_v0_2.tla` with an inline comment block recording the v0.5.1 property rewrite; the `cfg` file is unchanged. TLC has not been run; `TLC_STATUS_v0_5.md` continues to apply. Verification: programmatic substring check confirms `SubStar` or `EffectiveRevokedBy` appears within each of the three rewritten property bodies.

**T2P (Track 2 runner workdir bug).** Fixed in v0.5.1. `run_track2_minimal_benchmark_v0_5_1.py` resolves the validator and shapes paths relative to `__file__` (the script's own directory) rather than relative to `workdir`. The `workdir` argument now controls only where synthetic workloads and validation reports are written, not where the validator and shapes live. Also adds optional `validator`, `shapes`, and `generator` arguments with sensible defaults. The script can be run from any working directory.

**D005v2 (sensitivity range incorrect).** Fixed in v0.5.1. Paper §4.6 Finding 2 previously stated the Out-of-Art.22-Scope share ranges between [4%, 52%]. This range is not derivable from any subset of the 25-case sample. v0.5.1 replaces it with two explicit constructions, both anchored on  $k=6, n=25$ : under the F2-relevant subset flip (B7, B9, B15, B17, B22) the range is [16%, 36%]; under the all-seven inclusive-BORDERLINE flip (B7, B9, B14, B15, B16, B17, B22) the range is [16%, 44%]. Both arithmetic constructions are documented in-line so future readers can reproduce. The Wilson 95% CI [11.5%, 43.4%] is unchanged. Verification: `grep "[4%, 52%]"` returns zero hits.

**D008v2 (Establish framing in Appendix F.4).** Fixed in v0.5.1. F.4 opening rewritten to “Would establish, conditional on (i) the transition system in `mcs_revocation_model_v0_2.tla` being a faithful model of a concrete revocation pipeline, and (ii) successful TLC model-checking of the LTL formulas in §F.3 against bounded instantiations: ... Neither (i) nor (ii) is satisfied at v0.5.1 release.” Also fixed: version-string drift across the paper. Replaced `mcs_case_sample_v0_2.csv` with `_v0_3.csv` (4 occurrences), `mcs_mapping_matrix_v0_2.csv` with `_v0_3.csv` (2 occurrences), `mcs_revocation_model_v0_1.tla` with `_v0_2.tla` (2 occurrences), `mcs_sme_profile_v0_1.ttl` with `_v0_2.ttl` (multiple occurrences), and `report_translator_v0_1.py` with `_v0_2.py`. Verification: targeted `grep` for the legacy version strings returns zero hits.

### Housekeeping addressed in v0.5.1

**D009v2 (borderline count mismatch).** Fixed. Paper §4.6 now reads “six cases with the strict all-caps BORDERLINE flag (B7, B9, B15, B16, B17, B22) plus one case (B14) with a Profile-B-conditional borderline flag, seven cases under the inclusive reading.”

**D011v2 (§6.3 five vs six classes; bare gov:EvidenceArtifact).** Fixed. Paper §6.3 now states the temporal extension defines six classes (was “five”) and references `gov:SupervisoryEvidenceArtifact` (was bare `gov:EvidenceArtifact`). Shape count updated from “eleven core” to “fourteen core” reflecting the BLOCKER 2 additions.

**RPS (reproduce script column-name bug).** Fixed. `reproduce_package_counts_v0_5_1.py` reads from the actual `status_flag` and `confidence_grade` columns rather than the non-existent `status` and `confidence` columns. Programmatic run produces non-empty distributions: `status {IP: 30, RC: 8, SD: 10}; confidence {H: 15, L: 8, M: 25}`.

**TLC1 (run\_tlc.sh assumes jar in CWD without download note).** Fixed. The script now contains a comment block listing the official TLA+ release page URL (<https://github.com/tlaplus/tlaplus/releases/>), recommends pinning v1.8.0 or later, and supports `TLA_TOOLS_JAR` environment variable override. The script does not auto-download the jar; reproducibility requires a recorded jar version, which an implicit pull would defeat.

**DUP (legacy paper file).** Addressed. `Janssen-2026-MCS-Agent-Ontology-Working-Paper.md` is not included in the v0.5.1 release zip. If the user re-uploads it, ignore.

### Items deliberately not fixed in v0.5.1

**T1K (Track 1 kappa execution).** Not fixed: an external second coder is required. The v0.5.1 package contains the blinded sample, the kappa-plus-PABAK script, and the README; it does not contain coder results. Track 1 execution remains pre-registered as a follow-up revision commitment.

**T2V (Track 2 publication-grade benchmark).** Not fixed: requires pySHACL, Apache Jena SHACL, and TopBraid SHACL API runs with engine-version, hardware, and graph-size parameters recorded. The v0.5.1 minimal benchmark is a smoke-test scaffold only and the README explicitly forbids citing it as evidence for the Track 2 p95 claim.

**LMR (legal mapping external review).** Not fixed: requires an independent EU data-protection-and-AI-law specialist. The v0.5.1 package contains the review protocol and the 48-row CSV template; it does not contain a completed review.

**TLA03 / VSD (TLC model-checking transcript).** Not fixed: requires a TLC run with documented jar version, configuration parameters, state-space size, and full transcript. The v0.5.1 package contains the model file with v0.5.1 transitive-closure property rewrites, the `cfg`, and the run script; it does not contain a TLC transcript. `TLC_STATUS_v0_5.md` continues to be the authoritative status note.

**DCM (deep-coding matrix is provisional reconstruction).** Not fixed structurally: the deep-coding matrix at v0.5.1 remains a provisional round-1 reconstruction from the paper tables, not independent source coding. The paper text continues to label it as scaffolding.

---

## v0.4 ERRATA backlog reconciliation

---

The v0.4 ERRATA listed several “implemented corrections” that were partial or absent in the paper source at v0.5. v0.5.1 status per claim:

1. **Reframed paper from delivered “minimum” to “candidate Minimum Conceptual Set” under explicit counting convention.** Paper-level reframing was applied. The “explicit counting convention” was internally inconsistent until v0.5.1 (D004v2). v0.5.1 status: fixed.
2. **Distinguished “23 conceptual slots / net-new classes” from raw “69 owl:Class”.** This is the v0.4 ERRATA item that produced the v0.5 D004v2 defect: the paper conflated “23 slots” with “23 net-new classes” while the framework specification arithmetic gave 16. v0.5.1 separates these into three figures (23 slots, 16 net-new classes under strict reuse-zero accounting, 69 owl:Class

- declarations). v0.5.1 status: fixed.
3. **Replaced “validated by SHACL” with “structurally checkable by SHACL”.** Applied at v0.4. v0.5.1 status: unchanged.
  4. **Reframed “materially better than ad-hoc practice” as “testably better, pending validation”.** Applied at v0.4. v0.5.1 status: unchanged.
  5. **Reframed validation pack as “validation design”, not completed validation evidence.** Applied at v0.4. v0.5.1 status: unchanged.
  6. **Reframed SME language from “exemption” to “proportionality / prose substitution”.** Labels and prose updated at v0.4 with legacy URIs retained. v0.5.1 status: unchanged.
  7. **Fixed SME profile reference from `gov:BayesianEvidenceDecayShape` to `gov:EvidenceDecayShape`.** Applied at v0.4. The remaining SME-side dangling references (D003v2) were not addressed at v0.4 and are now addressed at v0.5.1.
  8. **Added anti-gaming note for SME prose substitution.** Applied at v0.4. v0.5.1 status: unchanged.
  9. **Added `primary_nonqualifying_flag` and `conditional_nonqualifying_note` to case sample.** Applied at v0.4. v0.5.1 status: unchanged.
  10. **Added `evidence_mapping_note` to mapping matrix.** Applied at v0.4. v0.5.1 status: unchanged.
  11. **Updated TLA+ sketch to v0.2 with `EffectiveRevokedBy`.** Applied at v0.4. The four properties were not rewritten to use the transitive closure until v0.5.1 (TLA02). v0.5.1 status: fixed.
  12. **Added TLC `.cfg` skeleton.** Applied at v0.4. The `cfg` is a skeleton; no TLC transcript exists. v0.5.1 status: unchanged.
  13. **Updated report translator shape map.** Applied at v0.4. v0.5.1 status: unchanged.
  14. **Added artefact manifest and reproducibility count script.** Applied at v0.4. The reproduce script had a column-name bug (RPS) that was not addressed at v0.5 and is now addressed at v0.5.1.

---

### Authoritative status as of v0.5.1

---

- The release artefacts are internally consistent across paper, framework specification, vocabulary, profiles, SME profile, scenarios, and reproducibility scaffolding.
- Validation evidence is pre-registered, not completed: Track 1 inter-rater coder result not produced, Track 2 publication-grade benchmark not run, TLC transcript not produced, external legal review not performed.
- The package is preprint-quality: defensible enough to publish for review, not a peer-reviewed or standards-validated artefact.

## Appendix C: Artefact manifest v0.5.1

---

Source file: `mcs_artifact_manifest_v0_5_1.md`

# MCS v0.5.1 Generated Artefact Manifest

Generated: 2026-04-25 (v0.5.1 release)

DOI (v0.5.1): [10.5281/zenodo.19758441](https://doi.org/10.5281/zenodo.19758441) License: CC BY 4.0

## Authority hierarchy

Machine-readable files are authoritative for the ontology, SHACL profiles, CSV coding, and scripts. Markdown files are human-readable specifications and supporting documentation. PDF and DOCX renderings are derived publication outputs; where rendering and markdown source disagree, the markdown source wins.

## RDF/OWL vocabulary

File: `mcs_vocabulary.ttl`

Metric	Count
RDF triples	517
owl:Class declarations	69
owl:ObjectProperty declarations	16
owl:DatatypeProperty declarations	24

Counting interpretation: - The 23-item claim refers to conceptual slots (paper §2.5 working figure). - The 16-item figure is the net-new-class count under strict reuse-zero accounting (5 DPV reuses in slots 6 to 10 and 2 PROV-O reuses in slots 3 and 19 contribute zero net-new classes; 23 minus 7 equals 16). - The 69-item figure is the raw owl:Class count in the Turtle source, including subtypes, support classes, and named instances.

## SHACL profiles

File: `mcs_profiles_v1_2.ttl`

Metric	Count
RDF triples	427 (v0.5: 371; +56 from three v0.5.1 NodeShapes)
sh:NodeShape declarations	21 (v0.5: 18; +3 from v0.5.1)

NodeShapes (alphabetical): - AgentActionShape - BiometricLogShape - CrossProfileGovernanceShape - DecisionShape - DelegationGrantShape - DriftIndicatorShape - EvidenceDecayShape - EvidenceQualityShape - HumanInterventionShape (added v0.5.1) - LogEventSubclassGranularityShape (added v0.5.1) - MonitoringObservationShape - PolicyVersionShape - ProfileA\_Art22ScopeShape - ProfileA\_DrawsStronglyOnShape - ProfileB\_Art22ScopeShape - ProfileB\_DrawsStronglyOnShape - ProfileB\_MeaningfulInterventionShape - ProfileDeltaShape - RevocationEventShape (added v0.5.1) - RoleAssignmentShape - SchufaUpstreamShape

Architectural breakdown: 14 core temporal-extension shapes (including 3 v0.5.1 SME-supporting shapes) plus 7 profile-layer shapes.

## SME proportionality profile

File: `mcs_sme_profile_v0_2.ttl`

Metric	Count
RDF triples	182
sh:NodeShape declarations	3

NodeShapes: - ExemptionSuspensionShape - ProseSubstitutionShape - SMEExemptionMarker

External shape references in this file all resolve against `mcs_profiles_v1_2.ttl` after the v0.5.1 BLOCKER 2 fix. Legacy URI names retain "Exemption" for backward compatibility; labels and comments use proportionality terminology.

## Mapping matrix

File: `mcs_mapping_matrix_v0_3.csv`

Metric	Count
Rows / legal cells	48
Columns	14

Status distribution: IP 30, RC 8, SD 10. Confidence distribution: H 15, M 25, L 8.

Per legal column status distribution:

Legal column	SD	IP	RC	Total
AT Act Annex III	2	2	2	6

AI Act Annex III	3	3	4	8
AI Act Art. 12	5	3	0	8
AI Act Art. 14	0	8	0	8
GDPR Art. 22	0	4	4	8
GDPR Art. 28	2	4	2	8
NIS2 Art. 23	0	8	0	8

## Case sample

File: `mcs_case_sample_v0_3.csv`

Metric	Count
Case rows	25
Columns	24
Primary NonQualifyingADM cases	6
Strict all-caps BORDERLINE cases	6 (B7, B9, B15, B16, B17, B22)
Inclusive borderline cases	7 (six strict plus B14 Profile-B-conditional)

Finding bucket distribution: Solely 15, Out-of-Art.22-Scope 6, Assisted 3, Pending 1.

NonQualifyingADM Wilson 95% CI on  $k=6$ ,  $n=25$ : [11.5%, 43.4%]. F2-relevant flip range: [16%, 36%]. All-seven inclusive flip range: [16%, 44%].

## Deep-coding matrix

File: `mcs_deep_coding_matrix_v0_5.csv`

Metric	Count
Rows	125 (25 coding targets x 5 ontologies)
Columns	7

Status note: this matrix at v0.5.1 remains a provisional round-1 reconstruction from the paper tables, not independent source coding. The paper text continues to label it as scaffolding, not evidence.

## Track 1 blinded sample

File: `track1_blinded_package/track1_blinded_case_sample_seed42_v0_5_1.csv`

Metric	Count
Rows	25
Columns	10 (down from 20 in v0.5)

Columns retained: `blind_order`, `case_id`, `jurisdiction`, `case_name`, `date`, `brief_facts`, `step1_operative_provision`, `source_citation`, `ecli_or_url`, `coder_notes`.

Columns redacted in v0.5.1: `step1_outcome`, `step2_adm_substrate`, `step2_outcome`, `step3_intervention_properties_failed`, `step3_outcome`, `step4_borderline`, `step4_alternative_coding`, `step5_profileB_reclassification`, `primary_nonqualifying_flag`, `conditional_nonqualifying_note`. Coders cannot reverse-engineer the author's classification path.

## SHACL example pass and fail reports

Files: `shacl_examples/scenario_minimal_pass_expected_validation_report.ttl` and `..._fail_...report.ttl`.

Pass report: `sh:conforms true`; 0 results. Fail report: `sh:conforms false`; 17 results (was 11 in v0.5; the 6 additional violations come from the v0.5.1 `HumanInterventionShape` firing on the unpopulated `gov:HumanIntervention` instance).

The pass scenario was updated in v0.5.1 to declare all six conjunctive properties on its `gov:HumanIntervention` instance, required by the new `HumanInterventionShape`. The fail scenario was left unchanged; its pre-existing unpopulated `gov:HumanIntervention` now produces six additional shape violations in addition to the original eleven.

## TLA+ model

Files: `tla_model_checking/mcs_revocation_model_v0_2.tla` and `..._v0_2.cfg`.

The v0.5.1 release retains the v0.2 file name. v0.5.1 changes are recorded as inline comment blocks within the file: - Three properties (`CancelsQueuedWork`, `StopsInflightUnderOptIn`, `ArtefactsEventuallyStale`) rewritten to quantify over the transitive closure via `\|E a \|in revoked : SubStar(a, d)`. - `StalenessPersists` left direct on artefacts; transitive closure not relevant in that body. - `TransitiveRevocation` property unchanged; was already correct in v0.5.

TLC has not been run. `TLC_STATUS_v0_5.md` continues to be the authoritative status note. The `cfg` file is unchanged.

## Pending external evidence

---

The following remain pending external execution and are not in scope for v0.5.1: - TLC model-checking transcript. - Track 1 second-coder kappa-plus-PABAK results. - Track 2 publication-grade SHACL benchmark (pySHACL, Apache Jena SHACL, TopBraid SHACL API). - Independent legal review of all SD/IP/RC matrix classifications.

## Validation performed in this release

---

- All TTL files parse without error.
- NodeShape count verified at 21 in profiles, 3 in SME profile, 0 in vocabulary.
- All SME-side shape references resolve.
- Reproduce script runs cleanly; mapping-matrix distributions populate from `status_flag` and `confidence_grade` columns.
- SHACL minimal pass scenario validates with `conforms=true`; fail scenario validates with `conforms=false` and 17 results.
- TLA+ properties contain SubStar or EffectiveRevokedBy in the transitive-closure positions per the v0.5.1 rewrite.
- Paper has zero hits for legacy version strings (`v0_2.csv`, `v0_1.tla`, `sme_profile_v0_1.ttl`, `report_translator_v0_1.py`).
- Paper has zero hits for `[4%, 52%]`, `Establish`, `conditional on`, and `23 net new`.
- Scenarios file has 12 unique section numbers, no duplicates.

## Appendix D: Scenarios v0.5.1

---

Source file: `mcs_scenarios_v0_5_1.md`

# MCS Worked Scenarios

**Companion artefact** to *A Supervisory-Evidence Ontology for Agentic AI under EU Law: Minimum Conceptual Set and Temporal Extension* (Janssen, April 2026).

**Filename:** `mcs_scenarios_v0.2.md` **Purpose:** seven realistic worked scenarios that stress-test the MCS by populating its 23 classes from event streams. Referenced in §7.1 of the paper. **Licence:** CC BY 4.0 **Version:** v0.2

## §1. Scope

The scenarios stress-test the MCS by walking through realistic event streams from controllers operating under EU regulatory regimes. Each scenario does the same thing: enumerate the agent activity, populate the relevant MCS classes, run both Profile A and Profile B SHACL validation, surface any cross-profile delta as a governance artefact, and identify which formal sketch (F1-F16) covers the substantive analysis.

The scenarios are illustrative, not normative. A controller running an actual production deployment must adapt the patterns to their specific operational and legal context. The MCS's release-state claim is structural: every event class can be represented; every Profile is intended to produce a determinate structural-validation outcome; the cross-profile delta is recordable as provenance.

Three scenarios produce single-profile classifications (Scenarios 1-3). One scenario produces a profile divergence (Scenario 4). Three scenarios exercise temporal-governance and incident-response semantics (Scenarios 5-7). Together they cover the seventeen thematic slots of the MCS in operationally varied form.

The scenarios use synthetic data. Any resemblance to specific named controllers is by structural analogy only, not by reference to actual data subjects, decisions, or operations.

## §2. Scenario 1 – Retail banking credit decision

### Context

A mid-sized retail bank operating in the Netherlands offers consumer credit (revolving credit lines €500-€25,000) through a partly-automated underwriting pipeline. The pipeline includes a third-party credit-reference score and a rules-based affordability calculator. Final decisions are issued by an underwriting team operating under documented authority and discretion criteria.

### Event stream (illustrative trace)

```
T0      Customer submits credit application (€10,000 line, 24 months)
T0+10s  Agent retrieves credit-reference score from third-party processor
T0+15s  Agent applies internal affordability rules (DTI, liquidity, history)
T0+20s  Agent generates decision recommendation (approve/decline + rate)
T0+45s  Underwriting officer reviews case; decision context surfaced
T0+8min Underwriting officer issues final decision (approve at 8.9% APR)
T0+8m5s Decision communicated to customer
T0+8m6s Audit-trail artefact emitted to evidence store
```

### MCS class population

MCS class	Instance
gov:Agent	bank-internal underwriting agent (SoftwareAgent) plus underwriting officer (HumanAgent)
gov:LegalRole	bank as Controller plus credit-reference provider as Processor
gov:RoleAssignment	controller-processor relationship, time-indexed
gov:Activity	application processing, score retrieval, rule application, recommendation generation, officer review, decision issuance
gov:ToolInvocation	credit-reference score query (third-party API call)
gov:Decision	final-decision issuance (gov:AssistedDecisionWithMeaningfulHumanIntervention candidate)
gov:Score	credit-reference score with drawsStronglyOn predicate populated under both profiles
gov:HumanIntervention	officer review with all six properties recorded
dpv:DataSubject	applicant
dpv:PersonalData	application data, financial data
dpv:LegalBasis	Art. 6(1)(b) contractual necessity
dpv:Purpose	credit-decision processing
gov:LogEvent	gov:OperationMonitoringLogEvent for each Activity
gov:SupervisoryEvidenceArtifact	audit-trail emission with SystemLogSource reliability weight

## Profile A vs Profile B outcome

- **Profile A (AP-inspired permissive, quantitative):** drawsStronglyOn evaluated via flip-rate at  $\epsilon=0.10$ . The bank's internal modelling shows that with the credit-reference score alone (ignoring affordability rules), 35% of decisions flip when the score is perturbed by 10%.  $M=0.50$  threshold not met. Profile A classifies the decision as Assisted with meaningful intervention.
- **Profile B (CNIL/German-guidance-inspired stricter, qualitative):** drawsStronglyOn evaluated via three judicial indicia. The third-party credit-reference is a specialised scoring model (Indicium 1: true). The refusal-frequency-given-low-score in this bank's historical population is 0.78 (Indicium 2: false, threshold 0.90 not met). The decision has documented human intervention with all six properties (Indicium 3: hasDocumentedIntervention = true, fails the negative form). Profile B classifies the decision as Assisted with meaningful intervention.

Both profiles produce Assisted classification. No cross-profile delta surfaces.

## Formal sketches engaged

F1 (Decision  $\times$  Art. 22), F2 (Score  $\times$  SCHUFA upstream), F9 (Human intervention  $\times$  authority utilisation).

## Stress-test outcome

Scenario admits MCS class population without extension. Both profiles produce determinate classification. Cross-profile delta absent. Internal-consistency validation: pass.

---

## §3. Scenario 2 – Hospital triage assistant

---

### Context

A teaching hospital in Belgium operates an AI-assisted triage system in its emergency department. The system ingests patient vital signs, presenting complaint, and brief clinical history; produces a triage-level recommendation (categories 1-5 per the Manchester Triage System). Final triage assignment is made by a triage nurse with documented training and authority. The system is high-risk under AI Act Annex III (healthcare).

### Event stream (illustrative trace)

```
T0      Patient registered at ED reception
T0+30s  Vital signs entered into triage system
T0+45s  Agent ingests vitals, complaint, history
T0+60s  Agent produces triage-level recommendation (level 2: very urgent)
T0+90s  Triage nurse reviews recommendation against clinical judgement
T0+5min Triage nurse confirms level 2 assignment
T0+5min Patient routed to acute-care zone
T0+5m2s Audit artefact emitted with patient context redacted per Art. 9 governance
```

### MCS class population

MCS class	Instance
gov:Agent	triage-assistant SoftwareAgent plus triage nurse HumanAgent
gov:LegalRole	hospital as Controller plus Provider (AI Act Art. 3(3))
gov:RiskClassification	gov:HighRiskAI per AI Act Art. 6 plus Annex III healthcare
gov:IntendedPurpose	triage-level recommendation per Manchester Triage System
gov:SectoralRegime	gov:MDR (Rule 11 if classified as Software as Medical Device by IGJ)
gov:SupervisoryScope	gov:IGJ (Inspectie Gezondheidszorg en Jeugd)
gov:Activity	triage processing
gov:Decision	triage-level assignment (gov:AssistedDecisionWithMeaningfulHumanIntervention)
gov:HumanIntervention	nurse review with sector-calibrated threshold (5min for clinical complexity)
dpv:PersonalData	patient vitals, history (Art. 9 special category)
dpv:LegalBasis	Art. 9(2)(h) provision of healthcare
gov:LogEvent	gov:OperationMonitoringLogEvent plus gov:RiskIdentificationLogEvent for Annex III deployment
gov:MonitoringObservation	triage-level recommendation accuracy versus nurse-assigned level (post-market monitoring under Art. 72)

## Profile A vs Profile B outcome

- **Profile A:** Sector-calibrated interventionDuration threshold of 5 minutes (clinical complexity). Nurse review of 5min15s exceeds threshold. All six HumanIntervention properties documented. Profile A classifies as Assisted.
- **Profile B:** Same evaluation; Profile B's qualitative drawsStronglyOn does not engage because the triage-level recommendation is not a probability score in the SCHUFA sense. Profile B classifies

as Assisted.

Both profiles converge. The critical Art. 22 question (whether the triage produces a “legal or significant effect”) is answered in the negative: triage is a healthcare-routing activity, not a decision producing legal effects per Art. 22(1) qualifier. The case engages Annex III healthcare classification but not Art. 22.

### Formal sketches engaged

F9 (Human intervention with sector calibration), F10 (Risk classification × Annex III).

### Stress-test outcome

Scenario admits MCS class population. Sector calibration of interventionDuration threshold (5min for clinical complexity) populates correctly. Annex III healthcare classification handled. Cross-profile delta absent. Internal-consistency validation: pass.

## §4. Scenario 3 – Municipal benefits eligibility

### Context

A Dutch municipality administers means-tested social-assistance benefits under the Participatiewet (Awb administrative law context). An AI-assisted eligibility-assessment tool ingests applicant-provided data plus authoritatively-sourced data from BRP, UWV, and Belastingdienst APIs; produces an eligibility recommendation with calculated benefit amount. Caseworkers issue the formal decision (besluit) with documented authority and discretion under Awb Art. 3:2 (zorgvuldige voorbereiding) and Art. 3:46 (motivering).

### Event stream (illustrative trace)

```
T0 Applicant submits benefit application via municipal portal
T0+1day Agent retrieves data from BRP, UWV, Belastingdienst APIs
T0+1d1m Agent applies eligibility rules
T0+1d2m Agent generates eligibility recommendation plus calculation
T0+2day Caseworker reviews case file with documented authority
T0+2d3h Caseworker issues besluit (approve/deny) with motivering
T0+2d3h+5m Besluit communicated to applicant per Awb Art. 3:41
T0+2d3h+6m Audit artefact emitted with full evidence chain
```

### MCS class population

MCS class	Instance
gov:Agent	eligibility-assessment SoftwareAgent plus caseworker HumanAgent
gov:LegalRole	municipality as Controller plus Deployer of high-risk AI (AI Act Annex III public benefits)
gov:RiskClassification	gov:HighRiskAI per Annex III
gov:SectoralRegime	gov:AWB (Algemene wet bestuursrecht) plus Participatiewet
gov:SupervisoryScope	gov:AP (algorithm supervisor) plus gov:RDI (regulatory sandbox if applicable)
gov:Activity	eligibility-data retrieval, rule application, recommendation generation, caseworker review, besluit issuance
gov:Decision	besluit (gov:AssistedDecisionWithMeaningfulHumanIntervention if review meets criteria)
gov:HumanIntervention	caseworker review with Awb Art. 3:2 authority and Art. 3:46 motivering
dvp:PersonalData	BRP, UWV, Belastingdienst-sourced data
dvp:LegalBasis	Art. 6(1)(c) legal obligation under Participatiewet plus Awb
gov:LogEvent	gov:RiskIdentificationLogEvent (Annex III) plus gov:OperationMonitoringLogEvent
gov:NonQualifyingADM	empty (case engages Art. 22; not in NonQualifyingADM scope)

### Profile A vs Profile B outcome

- **Profile A:** Caseworker review of 3 hours with documented Awb Art. 3:2 authority and full Art. 3:46 motivering. All six HumanIntervention properties met. Profile A classifies as Assisted.
- **Profile B:** Same review evaluated under stricter test. Property 6 (authorityActuallyExercised) tested by population-level deviation rate: this caseworker’s deviation rate over the past 12 months is 14% (above 10% threshold). All six properties met under Profile B. Classifies as Assisted.

Both profiles converge. Critical observation: the municipality’s caseworker pool population deviation rate is the operational evidence for Property 6. A municipality with a sub-10% deviation rate would fail Profile B and the case would reclassify as Solely.

### Formal sketches engaged

F9 (Human intervention authority utilisation), F10 (Risk classification).

### Stress-test outcome

Scenario admits MCS class population. Awb context handled via gov:SectoralRegime. The Property 6 deviation-rate test produces a concrete operational measurement that the municipality can monitor. Cross-profile delta absent given current deviation rate; would surface if rate drops below 10%. Internal-consistency validation: pass with sensitivity to operational metric.

## §5. Scenario 4 – Platform delivery rider penalty

### Context

A delivery platform operating in multiple EU jurisdictions uses an automated fraud-detection system to penalise riders for suspected non-compliance (artificial trip lengthening, location spoofing, cancellation gaming). The system flags suspicious patterns; a remote review team in another jurisdiction reviews flags and issues penalties (account warning, temporary deactivation, permanent deactivation). The Uber/Ola Amsterdam case sequence (B9, B10, B11, B12, B16) is the empirical anchor.

### Event stream (illustrative trace)

```
T0      Rider completes trip; trip data ingested
T0+5min Agent applies fraud-detection model; trip flagged
T0+10min Flag added to remote review queue
T0+2hour Remote reviewer opens flag; trip data displayed
T0+2h45s Remote reviewer issues 7-day deactivation
T0+2h45s+10s Rider notified of deactivation by app
T0+2h45s+11s Audit artefact emitted with Krakow-office reviewer ID
```

### MCS class population

MCS class	Instance
gov:Agent	fraud-detection SoftwareAgent plus remote reviewer HumanAgent
gov:LegalRole	platform as Controller (multi-jurisdictional)
gov:Activity	fraud-detection scoring, flag generation, reviewer review, deactivation issuance
gov:Decision	deactivation decision (PROFILE-DEPENDENT)
gov:Score	fraud-probability score
gov:HumanIntervention	remote-reviewer review (45 seconds)
gov:RoleAssignment	reviewer role with documented Krakow-office authority
dvp:LegalBasis	Art. 6(1)(b) contractual necessity (disputed under Art. 22(2)(a))
gov:LogEvent	gov:RiskIdentificationLogEvent

### Profile A vs Profile B outcome – DIVERGENCE

- **Profile A (AP-inspired permissive):** Sector-calibrated interventionDuration threshold of 15 seconds for high-volume fraud-detection alerts. Reviewer’s 45-second review exceeds threshold. Properties 1-5 all met (authority documented, understanding of logic established, relevant data considered, margin of discretion present). Property 6 (authorityActuallyExercised) tested by population-level deviation rate: Krakow office deviation rate is 8% (below 10% threshold). Property 6 fails under Profile A’s strict reading; intervention is “symbolic” per Hof Amsterdam B10. Profile A classifies as Solely. **Note:** even the permissive profile reaches Solely here because of the population-level metric.
- **Profile B (CNIL/German-guidance-inspired stricter):** Same evaluation; Property 6 fails identically. Profile B classifies as Solely. Additionally, qualitative drawsStronglyOn engages: fraud-probability score is a specialised scoring model (Indicium 1: true), so the score itself is in scope under Profile B.

Both profiles classify the decision as Solely, but for different reasons: - Profile A: HumanIntervention Property 6 fails due to population-level deviation rate. - Profile B: drawsStronglyOn engages additionally via qualitative indicium.

The cross-profile delta surfaces structurally: Profile A reaches Solely via the intervention path (Krakow review symbolic); Profile B reaches Solely via two paths (intervention symbolic plus score-itself-in-scope).

### Cross-profile delta artefact

```
gov:ProfileDelta a gov:CrossProfileGovernanceShape ;
gov:profileA gov:Solely_via_intervention_path ;
gov:profileB gov:Solely_via_intervention_AND_score_paths ;
gov:deltaSubstantive true ;
gov:requiresJurisdictionalRouting true ;
gov:complianceImplication "Score itself in Art. 22 scope under Profile B; pre-screening evidence retention obligations expand."@en .
```

## Formal sketches engaged

F1 (Decision × Art. 22), F2 (Score × SCHUFA upstream, profile-divergent), F9 (Human intervention × authority utilisation, sector-calibrated, deviation-rate failure).

## Stress-test outcome

Scenario admits MCS class population. Sector calibration of interventionDuration (15s for fraud detection) populates correctly. Population-level deviation-rate test for Property 6 produces concrete operational measurement that surfaces the symbolic-review failure mode (Hof Amsterdam B10 anchor). Cross-profile delta surfaces structurally even when both profiles classify identically: the delta is in the path, not just the outcome. Internal-consistency validation: pass with delta surfacing as designed.

---

## §6. Scenario 5 – Delegation revocation propagation

---

### Context

A SaaS B2B integrator operates an automated procurement-decision pipeline on behalf of a controller. The integrator has subdelegated decision-making authority to a sub-processor for a specific category of low-value transactions. The controller revokes the integrator's primary delegation (e.g., due to security incident or contractual termination). The MCS must represent transitive revocation per F6 sketch.

### Event stream (illustrative trace)

```
T0      Controller-Processor relationship established (DelegationGrant_d1)
T0+1week Processor-Subprocessor sub-delegation (DelegationGrant_d2 isSubdelegationOf d1)
T0+30day Routine procurement decisions issued under d2
T1      Controller revokes d1 (RevocationEvent_r1)
T1+5sec d1 propagation: BlocksNewWork applies to d1
T1+1sec d1 propagation: MarksArtifactsStale applies to d1 outputs
T1+1sec Transitive revocation: d2 revoked (RevocationEvent_r2 derived from r1)
T1+2sec d2 propagation: BlocksNewWork applies to d2
T1+2sec d2 propagation: MarksArtifactsStale applies to d2 outputs
T1+10sec Controller's audit-trail artefact emitted with full revocation chain
```

### MCS class population

MCS class	Instance
gov:DelegationGrant	d1 (controller-processor), d2 (processor-subprocessor)
gov:isSubdelegationOf	d2 isSubdelegationOf d1
gov:RevocationEvent	r1 (revokes d1), r2 (transitive, revokes d2)
gov:BlocksNewWork	applied to d1 and transitively to d2
gov:MarksArtifactsStale	applied to d1 outputs and transitively to d2 outputs
gov:LegalRole	controller, processor, subprocessor
gov:executedUnderRevokedAuthority	false for all decisions issued before T1
gov:cachedUnderRevokedAuthority	true for artefacts produced before T1 but accessed after
gov:LogEvent	gov:RiskIdentificationLogEvent for r1 and r2

### Profile A vs Profile B outcome

The revocation scenario does not engage Art. 22 directly (decisions issued before revocation remain authoritative). Both profiles handle the revocation propagation identically through the F6 sketch. Profile-invariant.

## Formal sketches engaged

F6 (Delegation × Art. 28 processor chain), with explicit transitive-revocation property. Linked to LTL formalisation in Appendix F of paper ([mcs\\_revocation\\_model\\_v0\\_1.tla](#) companion artefact).

## Stress-test outcome

Scenario admits MCS class population. Transitive revocation from d1 to d2 propagates within bounded latency ( $k \approx 1$  second per Appendix F.3.1). The five revocation effects (BlocksNewWork, MarksArtifactsStale as defaults; CancelsQueuedWork, StopsInflightWork, RequiresReview as opt-ins) populate correctly. Internal-consistency validation: pass; LTL semantics (Appendix F formulas F.3.1 transitive, F.3.2 BlocksNewWork persistence, F.3.5 staleness persistence) hold structurally.

---

## §7. Scenario 6 – Cross-border one-stop-shop

---

### Context

A pan-European online retailer (lead supervisory authority: Irish DPC under GDPR Art. 56) operates a personalised pricing engine across 14 EU member states. A French data subject contests Profile B (CNIL/German-guidance-inspired stricter) classification of pricing decisions made on French data. The cross-border cooperation procedure under GDPR Art. 60 engages: lead authority IE-DPC, concerned authority CNIL, plus the EDPB Art. 65 binding-decision pathway if disagreement persists.

## Event stream (illustrative trace)

```
T0      French data subject lodges complaint with CNIL
T0+5d   CNIL forwards to IE-DPC as lead authority per Art. 56
T0+30d  IE-DPC requests retailer's MCS evidence for the data subject's pricing decisions
T0+45d  Retailer emits MCS-conformant Turtle to IE-DPC hypothetical SSII endpoint (research-direction
in §8 item 2)
T0+60d  IE-DPC validates against Profile A (AP-inspired permissive aligned); finds Assisted
classification with all six HumanIntervention properties met
T0+75d  IE-DPC shares evidence with CNIL per Art. 60 cooperation
T0+90d  CNIL re-validates against Profile B (CNIL/German-guidance-inspired stricter aligned); finds
Solely under qualitative drawsStronglyOn (specialised scoring model indicium triggered)
T0+120d Cross-profile delta surfaces; Art. 60(3) consensus not reached
T0+150d Matter referred to EDPB under Art. 65(1) (a)
T0+10mo EDPB binding decision (illustrative outcome; actual outcomes vary)
```

## MCS class population

MCS class	Instance
gov:Agent	personalised-pricing SoftwareAgent
gov:LegalRole	retailer as Controller plus IE-DPC as lead supervisor plus CNIL as concerned supervisor plus EDPB
gov:SupervisoryScope	gov:IE_DPC (lead) plus gov:CNIL (concerned) plus gov:EDPB
gov:Decision	pricing decision (PROFILE-DIVERGENT classification)
gov:Score	personalised-price modifier (gov:usesSpecialisedScoringModel = true)
gov:HumanIntervention	absent (fully automated pricing)
gov:ProfileDelta	surfaces between Profile A and Profile B
gov:CrossProfileGovernanceShape	flags requiresJurisdictionalRouting = true

## Profile A vs Profile B outcome — DIVERGENCE

- **Profile A (AP-inspired permissive):** drawsStronglyOn evaluated quantitatively. Flip-rate at  $\epsilon=0.10$  produces  $M=0.42$  (below 0.50 threshold). Profile A classifies as not in Art. 22(1) scope under quantitative test. With no human intervention, the decision is gov:NonAutomatedDecision under Profile A's reading because the threshold is not met. (Note: this is a quirk of the quantitative test; the AP guidance on which Profile A is based reads Art. 22 narrowly.)
- **Profile B (CNIL/German-guidance-inspired stricter):** drawsStronglyOn evaluated qualitatively. Specialised scoring model present (Indicium 1: true). drawsStronglyOn predicate engages. Decision classified Solely.

Cross-profile delta surfaces. The delta is structural: Profile A says the decision is outside Art. 22 scope altogether; Profile B says it is squarely within and Solely automated. This is the maximum-divergence case.

## Cross-profile delta artefact

```
gov:ProfileDelta a gov:CrossProfileGovernanceShape ;
  gov:profileA gov:NotInArt22Scope ;
  gov:profileB gov:Solely ;
  gov:deltaSubstantive true ;
  gov:divergenceMagnitude "maximum"@en ;
  gov:requiresJurisdictionalRouting true ;
  gov:art60CooperationStatus "active"@en ;
  gov:art65EngagementProspective true ;
  gov:complianceImplication ""Profile divergence between IE-DPC (Profile A) and CNIL
  (Profile B) cannot be resolved by technical means. The delta IS the governance
  artefact. Art. 60 cooperation, Art. 63 consistency mechanism, or Art. 65
  binding decision are the legal-political resolution paths.""@en .
```

## Formal sketches engaged

F1 (Decision × Art. 22), F2 (Score × SCHUFA upstream, fully profile-divergent).

## Stress-test outcome

Scenario admits MCS class population. The maximum cross-profile divergence (NotInScope vs Solely) is representable as gov:ProfileDelta with jurisdictional-routing flag. The structural representation of divergence does not resolve the underlying legal disagreement; it makes the disagreement auditable, recordable, and transmissible across the Art. 60 cooperation procedure. The MCS's claim is not that it resolves IE-CNIL doctrinal divergence (no instrument can); the claim is that the divergence is no longer a hidden compliance liability but an explicit governance artefact. Internal-consistency validation: pass with structural divergence surfaced as designed.

---

## §8. Scenario 7 — NIS2 plus GDPR concurrent incident

### Context

A Dutch energy utility (Essential entity under NIS2 Art. 3, Annex I) operates an automated grid-balancing AI system. A misconfiguration causes the system to issue erroneous load-shedding commands

during a 2-hour window, affecting approximately 40,000 households (severe operational disruption per NIS2 Art. 23(3)) plus producing erroneous customer-facing decisions (potential GDPR Art. 22 engagement). NIS2 reporting cadence and GDPR breach-notification cadence both engage. The MCS must represent the concurrent incident plus the F4 (NIS2 significant incident) and F11 (delegation revocation × NIS2 incident window) sketches together with GDPR incident notification under Art. 33.

### Event stream (illustrative trace)

```

T0      Misconfiguration deployed via routine policy version change pv2 (supersedes pv1)
T0+15min Erroneous load-shedding commands issued; SignificantIncident detected
T0+15m1s NIS2 Art. 23 reporting clock starts; GDPR Art. 33 reporting clock starts
T0+30m  Operations team initiates rollback to pv1
T0+45m  Pv2 effectively revoked; pv1 restored as active policy
T0+1h   Incident contained; load-shedding stopped
T0+24h  NIS2 Art. 23(4) early-warning report submitted to NCSC
T0+72h  NIS2 Art. 23(4) full-incident report submitted to NCSC
T0+72h  GDPR Art. 33 breach notification submitted to AP
T0+1mo  NIS2 Art. 23(4) final report submitted; GDPR Art. 34 affected-subject notifications issued
        for the subset where Art. 22 engagement is established

```

### MCS class population

MCS class	Instance
gov:Agent	grid-balancing SoftwareAgent
gov:LegalRole	utility as Controller plus Operator of essential service (NIS2)
gov:SectoralRegime	gov:NIS2 (energy sector) plus GDPR Art. 22 for affected customer decisions
gov:SupervisoryScope	gov:NCSC (NIS2 incident) plus gov:AP (GDPR breach) plus gov:ACM_Energy (sectoral)
gov:Incident	the 2-hour load-shedding event
gov:detectionTime	T0+15min
gov:incidentWindow	[T0+15min, T0+1h]
gov:PolicyVersion	pv1 (pre-incident, pv2 (incident-causing), pv1 again (post-rollback))
gov:supersedesPolicy	pv2 supersedes pv1 (T0); pv1 supersedes pv2 (T0+45m) effectively
gov:Decision	erroneous load-shedding commands plus customer-facing decisions during incident window
gov:LogEvent	gov:IncidentLogEvent (NIS2 Art. 23) plus gov:OperationMonitoringLogEvent plus gov:RiskIdentificationLogEvent
gov:MemoryWrite	preserved under WORM semantics during incident window per F12
gov:DelegationGrant	grid-balancing authority delegation (default StopsInflightWork during incident per F11)
dpv:DataSubject	affected customers
gov:NonQualifyingADM	empty for the load-shedding decisions (operational, not Art. 22); populated for any customer-facing decisions decided on Art. 5 fairness or Art. 6 lawful basis if those grounds emerge

### Profile A vs Profile B outcome

Profile-invariant for the load-shedding decisions (operational/safety, not Art. 22). For customer-facing decisions issued during the incident window where Art. 22 may engage, the standard profile evaluation applies (per Scenarios 1–4).

### Formal sketches engaged

F4 (Tool invocation × NIS2 significant incident — load-shedding commands as ToolInvocation contributing to SignificantIncident), F11 (Delegation revocation × NIS2 incident window — pv2 revocation falls within incident window, default revocation behaviour shifts to StopsInflightWork), F12 (Memory write × NIS2 Art. 23(4) reporting cadence — incident-window writes preserved with WORM semantics for staged reporting), F13 (Policy version change × NIS2 Art. 21 measures — pv1→pv2→pv1 transitions recorded with measure deltas).

### Stress-test outcome

Scenario admits MCS class population for both NIS2 and GDPR concurrent regimes. The four NIS2-anchor sketches (F4, F11, F12, F13) all engage simultaneously and produce consistent constraints. The staged reporting cadence (24h early warning, 72h notification, 1 month final report) populates as required. The policy-version revocation propagation (pv2 effectively revoked at T0+45m) interacts correctly with the incident-window preservation rules. Cross-regime concurrent reporting (NIS2 plus GDPR Art. 33) produces no structural conflict; the same evidence base supports both. Internal-consistency validation: pass for the concurrent-regime stress test.

## §9. Scenario 8 — Hostile vendor SaaS control-plane telemetry refusal

## Context

A controller deploys a vendor-hosted agentic AI service for customer-complaint triage. The controller is the GDPR controller and AI Act deployer. The SaaS vendor provides the agent runtime, memory layer, tool-routing logic, and audit dashboard. After a contested automated escalation, the controller asks for raw event evidence. The vendor offers a PDF export but refuses to expose underlying RDF/JSON event streams, claiming the control-plane telemetry is proprietary.

### Event stream (illustrative trace)

```
T0      Customer complaint submitted
T0+3s   Vendor agent reads historical complaints from memory
T0+6s   Vendor agent calls sentiment-risk tool
T0+8s   Agent escalates case to debt-recovery workflow
T0+10s  Customer receives adverse communication
T0+1d   Controller requests full event evidence from vendor
T0+2d   Vendor supplies dashboard PDF only; raw tool/memory trace withheld
```

### MCS class population

MCS class	Instance
gov:Agent	vendor-hosted SoftwareAgent plus controller operations team
gov:LegalRole	controller as Controller and Deployer; vendor as Processor and possible Provider
gov:RoleAssignment	controller-vendor role attribution with time interval
gov:ToolInvocation	sentiment-risk tool call, routing API call
gov:MemoryEvent	historical complaint memory read
gov:Decision	escalation to debt-recovery workflow; Art. 22 significance requires legal review
gov:SupervisoryEvidenceArtifact	vendor PDF export plus missing raw telemetry artefact
gov:ProvenanceChain	incomplete chain because raw event stream is withheld
gov:PolicyVersion	vendor routing-policy version if disclosed; absent if proprietary
gov:NonQualifyingADM	not populated unless the case is decided on Arts 5, 6, 9, 25, or 35 rather than Art. 22

### Profile A vs Profile B outcome

- **Profile A:** May pass structural classification only if the dashboard export contains enough quantitative fields to populate the required shapes.
- **Profile B:** Likely fails or produces a governance delta if vendor evidence does not show meaningful intervention, score reliance, and raw event provenance.

### Formal sketches engaged

F3 (tool-chain propagation), F5 (memory read), F8 (evidence quality), F11/F12 (incident/evidence preservation where applicable).

### Stress-test outcome

This scenario is adversarial because the vendor can satisfy narrative compliance while withholding machine-readable evidence. The MCS does not solve the contractual-access problem. It surfaces it. The required control is procurement-level: controller contracts must require exportable event evidence, shape-conformant telemetry, retention periods, and supervisor-access procedures. Without those contractual controls, structural conformance can be gamed.

---

## \$10. Scenario 9 — Public-sector high-risk procurement with foundation-model provider conflict

---

### Context

A public authority procures a case-prioritisation assistant built on a foundation-model API. The public authority is the deployer and GDPR controller for case data. The systems integrator configures prompts, retrieval sources, and workflow tooling. The foundation-model provider controls base-model updates and safety policies. A model update changes output behaviour and increases false-positive prioritisation for a protected group.

### Event stream (illustrative trace)

```
T0      Public authority approves production deployment
T0+30d  Foundation-model provider updates base model
T0+31d  Integrator changes retrieval prompt template
T0+32d  Drift indicator triggers on protected-group false-positive rate
T0+33d  Authority suspends automated prioritisation
T0+34d  Supervisor asks who controlled the relevant change
```

### MCS class population

MCS class	Instance
gov:Agent	public authority, integrator, foundation-model provider, software agent
gov:LegalRole	Controller, Processor, Deployer, Provider, possible distributor/importer roles
gov:RoleAssignment	time-indexed attribution for each party
gov:PolicyVersion	authority policy, integrator prompt policy, provider model-release policy
gov:RiskClassification	likely HighRiskAI if Annex III public-service access is engaged
gov:IntendedPurpose	case-prioritisation purpose as declared by deployer/provider
gov:MonitoringObservation	false-positive rate by group
gov:DriftIndicator	performance or control drift after provider update
gov:Incident	populated if threshold for serious/significant incident is met
gov:SupervisoryEvidenceArtifact	model-release note, prompt-change log, monitoring report, suspension record

### Profile A vs Profile B outcome

The Article 22 classification depends on whether prioritisation produces legal or similarly significant effects and whether human intervention is meaningful. The more important stress is role attribution: the MCS can represent multiple time-indexed legal roles, but it cannot by itself resolve whether the foundation-model provider, integrator, or public authority is responsible for the harmful change.

### Formal sketches engaged

F10 (Annex III risk classification), F11/F12/F13 (incident and monitoring), F15 (model/policy drift), F16 (role and evidence reconstruction, if added in later revisions).

### Stress-test outcome

Scenario admits MCS representation but exposes a governance gap: provider/deployer responsibility allocation must be supported by procurement terms, model-change notice duties, monitoring thresholds, rollback rights, and evidence-access clauses. The MCS is a representational layer, not a liability resolver.

## §11. Cross-scenario observations

Three observations follow from the seven scenarios.

**Observation 1: Profile divergence is the exception, not the rule.** Five of seven scenarios produce profile-invariant outcomes (Scenarios 1, 2, 3, 5, 7). Scenario 4 (platform delivery) is the ordinary divergence case: a substantive cross-profile delta in the classification path that the dual-profile architecture was designed to surface. Scenario 6 (cross-border one-stop-shop) is an intentionally extreme cross-border stress case: by construction, the AP-permissive and CNIL-strict readings produce a maximum divergence in classification outcome (NotInScope vs Solely) on the same underlying decision. Scenario 6 stress-tests the limit case rather than representing typical deployment. Most production agentic AI deployments will run under a single profile for most decisions.

**Observation 2: Sector calibration is operationally consequential.** Scenarios 2, 3, and 4 use different interventionDuration thresholds (5 minutes for clinical triage, no override for benefits casework, 15 seconds for fraud detection). Each calibration is declared per case rather than applied silently, and each produces a different operational test. A controller running multiple calibrations in different products needs the calibrations to be visible in the supervisory evidence rather than embedded in opaque thresholds.

**Observation 3: The four Profile-B HumanIntervention properties cluster in two paths.** Properties 1-4 (hasAuthorityToDeviate, hasUnderstandingOfLogic, consideredAllRelevantData, hasMarginOfDiscretion) are ex-ante design properties about the reviewer's role; they are typically met or not met based on documented role definition. Properties 5-6 (interventionDurationSeconds, authorityActuallyExercised) are ex-post operational properties about what the reviewer actually did and how the population behaves. Profile divergence almost always concentrates in Properties 5 and 6 rather than 1-4, because the design properties are knowable in advance while the operational properties are emergent. This is the empirical pattern to monitor in the validation pack Track 1 work.

## §12. Limits

The seven scenarios are illustrative, not exhaustive. They do not cover:

- **Pure inference-time agent operations** without state-bearing decisions (e.g., conversational assistants without persistent customer-facing impact). The MCS likely under-fits these by design; the gap is intentional.
- **Multi-tenant SaaS-on-SaaS topologies** with deep delegation chains beyond two levels. F6 sketch handles two-level chains; deeper chains require empirical evaluation.
- **AI-system-to-AI-system delegation** (one autonomous system delegating to another). This is

research-conjecture territory; the MCS does not commit to a representation here. See Section 8 research-agenda items for further work.

- **Real-time safety-critical systems** (autonomous vehicles, surgical robotics, industrial control beyond grid balancing) where the millisecond-scale operational constraints exceed LTL's expressive power. TCTL or real-time logics would be required.
- **Cross-jurisdictional supply-chain shadow agents** where the deployer-provider boundary blurs across multiple legal regimes. The MCS handles two-jurisdiction cases (Scenario 6) but more complex chains require empirical evaluation.

The seven scenarios cover the seventeen thematic slots of the MCS in operationally varied form. They do not cover every conceivable agentic AI deployment pattern. Track 3 of the validation pack (`validation_pack_v1.md` §5) addresses contributor topologies that fall outside these scenarios.

---

*End of `mcs_scenarios_v0_2.md`. © 2026 J. Janssen, Apparens. CC BY 4.0.*

---

## Appendix E: Coding rubrics v0.2

---

Source file: `coding_rubrics_v0_2.md`

# MCS Coding Rubrics

**Companion artefact** to *A Supervisory-Evidence Ontology for Agentic AI under EU Law: Minimum Conceptual Set and Temporal Extension* (Janssen, April 2026).

**Filename:** `coding_rubrics_v0_2.md` **Purpose:** enable independent reproduction of the five coded artefacts in the paper. Track 1 of the validation pack (Appendix D, Section 8 item 14) commissions an independent second coder against this document. **Licence:** CC BY 4.0 **Version:** v0.2 **Coder of record (round 1):** J. Janssen, 23 April 2026.

## §1. Scope and structure

The paper produces five coded artefacts:

1. **Survey coverage matrix** (§4.1, Tables 4.1a–4.1d). 14 ontologies × 18 concepts = 252 cells, three-value rating present/partial/absent.
2. **Deep-coding reusability matrix** (Appendix F.3 reference). 25 MCS class slots × 5 ontologies (PROV-O, GDPRov, DPV, Agent Spec, ACP Framework) = 125 cells, three-value rating reusable/partially reusable/not reusable under owl:imports plus operational-criteria preservation.
3. **Mapping matrix with status flags** (Appendix A, full 48-cell matrix in `mcs_mapping_matrix_v0_2.csv`). 8 agent-behaviour primitives × 6 legal categories = 48 cells, status flag {SD, IP, RC} plus confidence {H, M, L}.
4. **Case classification** (Appendix B, full record in `mcs_case_sample_v0_2.csv`). 25 EU ADM enforcement cases, five-step protocol.
5. **Borderline-case sensitivity** (Appendix B.5, sensitivity flags in `mcs_case_sample_v0_2.csv`).

Each rubric below is presented in a uniform structure: scope, materials access, decision rule, output schema, tie-breaker, randomisation, reliability target.

## §2. Survey coverage matrix prompt (Tables 4.1a–4.1d)

### Scope

For each (ontology, concept) cell: assign one of three values.

- **Present (●):** A formal class or property in a published schema matches the concept under synonymous or strictly narrower terms. Evidence required: an OWL class declaration, an RDF property, a SHACL shape, or an explicit named construct in schema documentation.
- **Partial (P):** The concept is covered conceptually in prose documentation but lacks a formal construct, OR the formal construct is identifiable-but-imperfectly-overlapping (e.g., DPV-RISK extension covers Risk classification via a separate module rather than DPV core).
- **Absent (○):** No construct or concept matching the row label is returned by structured search of the ontology’s published schema and prose documentation.

### Materials access

The coder must have access to: - Each ontology’s canonical specification document (schema file plus accompanying prose). - The 18 concept definitions specified in §4.1 of the paper.

Canonical sources for the 14 ontologies in scope: - Aviso, Atlan, Microsoft Fabric IQ, Skan: vendor documentation pages plus public schema where available. - Agent Spec (Open Agent Specification): GitHub repository plus published specification. - ACP Framework (Agent Control Protocol): published specification. - OntoBoom, Foundation AgenticOS: vendor documentation. - PROV-O: W3C Recommendation 30 April 2013. - OWL-Time: W3C Recommendation 2017 plus 2022 Candidate Recommendation update. - GDPRov: Pandit and Lewis SEMANTiCS 2017 plus published vocabulary. - DPV: W3C DPVCG Final Community Group Report v2.3 25 February 2026. - ISO/IEC 42001 + 23894 + 5338 (grouped): published prose standards. - CEN-CENELEC JTC 21 drafts (prEN 18286, 18228, 18284): drafts available via national standards bodies subject to confidentiality where applicable.

### Decision rule

For each cell, in this order:

1. Search the ontology’s machine-readable schema for a construct matching the row concept under synonymous or strictly narrower terms. If found: code Present (●).
2. Search the ontology’s prose documentation for a recognisable named concept or descriptive prose matching the row concept. If found without formal construct: code Partial (P). If formal construct exists in prose-only standards (ISO/IEC, JTC 21), code Partial (P) under the strict-legend reading rather than Present.
3. If neither search returns a match: code Absent (○).

### Output schema

For each cell: cell value plus optional 1–2 sentence justification noting the formal construct name (if Present) or the prose source (if Partial). Example:

DPV × Personal-data category: ● (formal construct `dpv:PersonalData` with subclass hierarchy in DPV)

```
core).
ISO/IEC × Risk classification: P (prose-level coverage in ISO/IEC 23894 risk-management terminology;
no formal class).
```

### Tie-breaker

If a concept is covered by an extension module rather than core (e.g., DPV-RISK extension for Risk classification), code as Partial with parenthetical note specifying the extension module. The strict-legend reading prefers core-coverage for the Present rating.

### Randomisation

Round 1: cells coded in primitive-major order (concept-by-concept across ontologies). Round 2: cells coded in random order with seed `np.random.seed(42)` (ontologies and concepts shuffled independently).

### Reliability target

Track 1 inter-rater Cohen's  $\kappa \geq 0.70$  (substantial agreement per Landis and Koch 1977). Failure at  $\kappa < 0.60$  triggers rubric revision rather than cell-level reconciliation.

---

## §3. Deep-coding reusability matrix prompt (Appendix F.3)

---

### Scope

For each (MCS class slot, ontology) cell: assign one of three values.

- **Reusable:** The MCS class can be operationalised by direct `owl:imports` of the ontology's existing class with no operational-criteria loss. Five MCS slots are direct DPV reuse under this convention (Slots 6–10): DataSubject, PersonalDataCategory, PurposeOfProcessing, LawfulBasis, ConsentRecord.
- **Partially reusable:** The MCS class can use the ontology's class as a parent or subclass-of target, but adding operational properties is required to preserve supervisory-evidence semantics (e.g., `gov:Activity` subclassOf `prov:Activity` adds AI Act-specific properties).
- **Not reusable:** The MCS class either has no analogue in the ontology, or the ontology's analogue lacks the operational properties required for supervisory ingestion.

### Materials access

Same materials as §2 plus the §4.3 MCS class definitions (Tables 4.3a, 4.3b).

### Decision rule

For each cell:

1. Find the ontology's closest analogue to the MCS class under either name match or definitional match.
2. Test whether the MCS class can be defined via `rdfs:subClassOf` of the analogue without losing operational properties specified in the MCS class definition.
3. Code Reusable if zero new properties needed; Partially reusable if subclass-of relationship preserved with added properties; Not reusable if no usable analogue exists.

### Output schema

Three-value rating per cell plus 1-sentence note specifying the analogue (if any) and the lost-or-added properties.

```
gov:HumanIntervention × DPV: not reusable (DPV has no MeaningfulIntervention construct; six-property
schema is MCS-side work).
gov:Activity × PROV-O: reusable (direct prov:Activity reuse with no new properties; counted as 0 net
new under §2.5 convention).
gov:DelegationGrant × ACP Framework: partially reusable (ACP's delegation construct provides
cryptographic guarantee; MCS adds gdpr-arts-28 plus ai-act-art-25 anchoring).
```

### Tie-breaker

Where an ontology has multiple candidate analogues, prefer the most specific match. Example: `gov:DelegationGrant` maps to `prov:Delegation` (more specific) rather than `prov:Attribution` (more abstract).

### Randomisation

Same as §2.

### Reliability target

Same as §2 ( $\kappa \geq 0.70$  substantial agreement;  $\kappa < 0.60$  triggers rubric revision).

---

## §4. Mapping matrix status and confidence prompt (Appendix A;

`mcs_mapping_matrix_v0_2.csv`)

---

## Scope

For each (agent-behaviour primitive, legal category) cell: assign one status flag and one confidence grade. 48 cells total under Profile A; the same 48 may be re-coded under Profile B if Profile B differs, recorded in a separate `profile_B_status` column.

## Status flags

- **[SD] Statutory derivation:** The classification rule follows directly from statutory text, recital, or settled CJEU jurisprudence. Test: a competent EU lawyer reading the operative statutory provision would derive the rule without exercising interpretive judgement on contested questions.
- **[IP] Interpretive proposal:** The rule is a defensible reading of the source texts but rests on an interpretive step on which competent authorities or courts could differ. Test: a competent EU lawyer could disagree with the rule on a defensible alternative reading without committing legal error.
- **[RC] Research conjecture:** The rule extends source materials beyond what they address. No statutory provision, recital, or jurisprudence directly grounds it. Test: the rule would require either new legislation, new case-law, or new DPA guidance to become a settled mapping.

## Confidence grades

- **H (high):** The rule is anchored in current statutory text or jurisprudence with high doctrinal stability; controllers can rely on it.
- **M (medium):** The rule is defensible but rests on synthesis or extension; controllers should be aware that supervisors may apply alternative readings.
- **L (low):** The rule is research-stage; controllers should not rely on it as compliance basis.

## Materials access

The coder must have access to: - The 8 primitives and 6 legal categories specified in Appendix A.2. - Source provisions per legal category as listed in Appendix A.3. - The formal sketches F1-F16 (Appendix A.4-A.19; Section 5.3 inline sketches F1, F2, F4). - WP251rev.01 plus AP handvatten (July 2025).

## Decision rule

For each cell:

1. Find the operative statutory provision in the legal-category column.
2. Determine whether the primitive is named in the provision or only inferable via interpretation.
  - Named: candidate SD.
  - Inferable: candidate IP.
  - Not inferable: candidate RC.
3. Apply confidence grade based on doctrinal stability of the source.

The 48 cells in `mcs_mapping_matrix_v0_2.csv` provide canonical round-1 codings as reference. Round-2 self-consistency under shuffled cell order tests within-coder stability.

## Output schema

Per cell: status flag {SD, IP, RC}, confidence grade {H, M, L}, formal-sketch reference (F1-F16 or n/a), source anchor (specific statutory paragraph or jurisprudential citation), one-sentence rule summary, plus rationale of 2-4 sentences.

The CSV columns of `mcs_mapping_matrix_v0_2.csv` formalise this output schema: `cell_id`, `row_primitive`, `col_legal_category`, `status_flag`, `confidence_grade`, `formal_sketch`, `profile`, `source_anchor`, `operative_provision`, `rule_summary`, `rationale`, `limits_and_edges`, `coder_notes`.

## Tie-breaker

Where two competent readings produce different status flags: - IP versus SD → default to IP (the more conservative reading). - RC versus IP → default to RC (the more conservative reading). The conservative-reading preference avoids overclaiming statutory derivation.

## Randomisation

Round 1: cells coded in row-major order (primitive-by-primitive across legal categories). Round 2: cells coded in random order with seed `np.random.seed(42)`.

## Reliability target

Status flag agreement:  $\kappa \geq 0.70$  substantial. Confidence grade agreement:  $\kappa \geq 0.60$  (reflects greater interpretive subjectivity in confidence than in status). Failure on either triggers rubric revision.

---

## §5. Profile A vs Profile B coding prompt

Under Profile B (CNIL-strict), several cells in the mapping matrix may shift status or confidence due to different operationalisation of the meaningful-intervention test and the `drawsStronglyOn` predicate. The Profile B re-coding is recorded as separate columns `profile_B_status` and `profile_B_confidence` in `mcs_mapping_matrix_v0_2.csv` if and when a coder runs the parallel pass.

## Scope of expected Profile B differences

- **C19 Delegation grant × Art. 22:** Profile B may upgrade IP to SD where the delegation includes Art. 22 decision authority and the meaningful-intervention test is applied strictly per Profile B's qualitative indicia.
- **C25 Delegation revocation × Art. 22:** similar parity with C19.
- **C37 Autonomous goal setting × Art. 22:** Profile B may upgrade RC to IP if national-court practice begins to apply the SCHUFA upstream-attribution path to goal-setting events.
- **Other cells:** most cells are profile-invariant; the matrix value remains unchanged.

## Decision rule

For each cell where Profile B may differ:

1. Apply the Profile B operationalisation: stricter intervention test (six conjunctive properties including authorityActuallyExercised), qualitative drawsStronglyOn (three judicial indicia per §6.4).
2. If the Profile B reading shifts the cell's status (typically by raising confidence due to stricter criteria producing more determinate classifications), record the shift.

## Output schema

Two new columns appended to `mcs_mapping_matrix_v0_2.csv` if the coder runs Profile B pass: `profile_B_status`, `profile_B_confidence`. Cells unchanged remain blank in the new columns.

## Reliability target

Profile B coverage is exploratory. No κ target; coder consistency reported as % of cells where Profile A and Profile B agree.

---

## §6. Case classification prompt (Appendix B; `mcs_case_sample_v0_2.csv`)

---

### Scope

For each of 25 cases: apply the five-step protocol from §2.7 and produce per-case Step 1-5 outputs plus primary classification, sensitivity flag, and NonQualifyingADM subclass mapping where applicable.

### Materials access

The coder must have access to: - The 25 case records (Table B.4 in the paper plus the brief facts and source citations in `mcs_case_sample_v0_2.csv`). - The five-step protocol definition (§2.7). - WP251rev.01 plus AP handvatten July 2025. - The six HumanIntervention property definitions (Table 4.3a Slot 12). - The NonQualifyingADM subclasses (Slot 23).

### Decision rule (five-step protocol)

**Step 1. Operative-provision test.** What GDPR or equivalent provision is the decision actually grounded on? If the operative ground is GDPR Arts 5, 6, 9, 25, or 35, or Art. 8 ECHR, rather than Art. 22: the case is classified Out-of-Art.22-Scope and populates the NonQualifyingADM structural category.

**Step 2. AI involvement test.** Was ADM or profiling present in the factual substrate? If no, the case is NonAutomated and exits the protocol.

**Step 3. WP251rev.01 meaningful-intervention test.** If ADM is present and Art. 22 is operative, does the human intervention in the record meet all six WP251rev.01 plus AP handvatten properties (hasAuthorityToDeviate, hasUnderstandingOfLogic, consideredAllRelevantData, hasMarginOfDiscretion, interventionDurationSeconds above threshold, authorityActuallyExercised)? If yes, Assisted. If no, Solely. Any property failure is recorded individually.

**Step 4. Borderline cases.** A case is borderline when Art. 22 reasoning is present in the decision but the operative ground is Arts. 5, 6, 9, 25, or 35. Primary coding per Step 1 (Out-of-Art.22-Scope); permissive alternative per Steps 2 and 3. Both recorded.

**Step 5. Profile B cross-check.** Under Profile B (CNIL-strict), the threshold for solely automated is lower, which may reclassify borderline Assisted cases as Solely. Any reclassification recorded.

### Operationalised thresholds

- **interventionDurationSeconds default:** 60 seconds. Defensible operational presumption based on AP handvatten "non-symbolic" requirement plus working assumption that symbolic review can be performed in under one minute per decision.
- **Sector calibration upward:** complex clinical triage decisions, threshold 5 minutes (300 seconds).
- **Sector calibration downward:** high-volume fraud-detection alerts, threshold 15 seconds, conditional on documented evidence that the reviewer examines case-specific features rather than rubber-stamping.
- The sector calibrations are declared per case in the `coder_notes` column, not applied silently.

### Output schema

CSV columns of `mcs_case_sample_v0_2.csv`: `case_id`, `jurisdiction`, `case_name`, `date`, `brief_facts`, `step1_operative_provision`, `step1_outcome`, `step2_adm_substrate`, `step2_outcome`, `step3_intervention_properties_failed`, `step3_outcome`, `step4_borderline`, `step4_alternative_coding`, `step5_profileB_reclassification`, `primary_classification`, `finding1_bucket`, `sensitivity_flag`, `grounds_NonQualifyingADM_subclass`, `supports_finding`, `source_citation`, `ecli_or_url`, `coder_notes`.

### Tie-breaker

Where Step 1 produces ambiguity over operative ground (e.g., a decision cites both Art. 5 and Art. 22): - Primary coding: the provision identified by the deciding authority as the operative ground. - Alternative coding: the other identified provision recorded in `step4_borderline=Yes` with `step4_alternative_coding` populated.

### Randomisation

Round 1: cases coded in B1-B25 case-id order. Round 2: cases coded in random order with seed `np.random.seed(42)` to test within-coder consistency.

### Reliability target

Track 1 (independent second coder): Cohen's  $\kappa \geq 0.70$  substantial agreement on primary classification;  $\kappa \geq 0.60$  on borderline-case sensitivity flags. Failure at  $\kappa < 0.60$  on primary classification triggers rubric revision rather than cell-level reconciliation.

---

## §7. Borderline-case sensitivity prompt

---

### Scope

For each of 25 cases: assign sensitivity flag {Stable, BORDERLINE, Stable post-appeal, Stable under Profile A; uncertain under Profile B, Borderline under Profile B, Pending, Stable for interpretive purposes}.

### Decision rule

A case carries a sensitivity flag of BORDERLINE if any of the following holds: 1. Step 4 of the case-classification protocol returns Yes (i.e., Art. 22 reasoning present but operative ground is non-Art. 22). 2. The case has been classified differently at first instance and on appeal (B16 reversed by B10 is the canonical example). 3. The Step 3 intervention-properties test produces a marginal result (one or two properties ambiguous in the record). 4. Profile B reclassification (Step 5) produces a different primary classification from Profile A.

A case carries the flag *Stable* if none of the four conditions hold.

Refinements: - *Stable post-appeal*: The case was BORDERLINE prior to appeal but the appeal authoritatively settled the classification (e.g., B10 stable post-appeal at Solely after first-instance Assisted). - *Stable under Profile A; uncertain under Profile B*: Profile A produces a determinate result but Profile B's stricter test could reach a different classification given partial evidence. - *Pending*: classification awaits a final supervisory determination (e.g., B21 AP Belastingdienst toezichtarrangement).

### Output schema

Single-column flag in `sensitivity_flag` column of `mcs_case_sample_v0_2.csv`. Rationale recorded inline in `coder_notes`.

### Materials access

Same as §6.

### Tie-breaker

Where two coders disagree on borderline-vs-stable assignment, the more conservative reading (BORDERLINE) is preferred. The default bias is towards over-flagging sensitivity, which produces wider sensitivity ranges in Findings 1-4 but reduces false-stability claims.

### Randomisation

Same as §6.

### Reliability target

$\kappa \geq 0.60$  (lower than primary classification because borderline assignment is a meta-judgement with greater interpretive variance).

---

## §8. Limits

---

The publication of these rubrics does not eliminate interpretive judgement. The following limits apply.

1. **Operationalised thresholds are defensible but not statutory.** The 60-second

interventionDuration default, the 0.10 Profile A flip-rate  $\epsilon$ , the 0.50 Profile A M threshold, the 0.90 Profile B refusal-frequency threshold, and the 6/12/24-month evidence-decay half-life priors are all controller-configurable and supervisor-rejectable. Coders applying the rubrics should record divergence from defaults in `coder_notes` rather than apply alternatives silently.

2. **The SD/IP/RC boundary is itself interpretive.** A cell that one competent lawyer reads as SD may read as IP to another. The conservative-reading preference (§4 tie-breaker) reduces but does not eliminate this variance.
3. **Sector calibration is judgement-laden.** The clinical-triage 5-minute threshold and fraud-detection 15-second threshold are illustrative, not statutory. Sector calibration in jurisdictions or domains not covered by the paper requires the coder to declare the calibration explicitly.
4. **Conflict-of-interest disclosures are self-reported.** The paper's commercial-interest disclosure (§1, §7.8) does not appear in the rubrics because it is an authorial-disclosure obligation, not a coder-side judgement. Coders running Track 1 should disclose any conflict of interest at submission.
5. **Round-2 audit is intra-rater not inter-rater.** Self-consistency testing against round 1 measures within-coder stability, not inter-subjective validity. Track 1 (independent second coder) is the inter-rater pass.
6. **The 25-case sample is selection-biased.** Cases reach courts or DPAs because they were contested. Population-level ADM distribution cannot be inferred from litigated-sample distribution. This caveat applies to any quantitative claim derived from the case-classification artefact.

---

## §9. Track 1 contributor protocol

---

### Eligibility

Independent coder must have: - LL.M. or equivalent in EU data protection law, or equivalent professional experience. - Working familiarity with WP251rev.01, SCHUFA, Dun and Bradstreet, and AP handvatten. - No prior involvement with the paper's authorship or review.

### Procedure

1. Receive the rubrics document plus blinded copies of `mcs_case_sample_v0_2.csv` (with `primary_classification`, `finding1_bucket`, `sensitivity_flag`, `grounds_NonQualifyingADM_subclass`, and `supports_finding` columns redacted) and `mcs_mapping_matrix_v0_2.csv` (with `status_flag` and `confidence_grade` columns redacted).
2. Apply the §2-§7 rubrics independently.
3. Submit completed CSVs.
4. Cohen's  $\kappa$  computed per artefact: case classification (primary), case classification (sensitivity flag), mapping matrix (status flag), mapping matrix (confidence grade), survey matrix (3-value rating), deep-coding matrix (3-value rating).
5. Per-cell disagreements catalogued with both codings retained for transparency.

### Failure handling

If  $\kappa < 0.60$  on any artefact: - Rubric revision rather than cell-level reconciliation. - Publication of the failure as part of the v2.0 commitment. - Re-coding by a third independent coder against revised rubrics.

If  $\kappa \geq 0.60$  but  $< 0.70$ : - Per-artefact  $\kappa$  reported. - Cells with disagreement noted in publication metadata.

If  $\kappa \geq 0.70$ : - Substantial agreement reported. - Original codings retained as canonical.

### Compensation

This validation programme operates on volunteer participation. There is no compensation, no honorarium, no advisory time exchange, and no other consideration offered to volunteers beyond co-citation in subsequent paper revisions where applicable. Volunteers retain full ownership of their submission data and may publish independently. See `validation_pack_v1.md` §2 for the rationale behind the volunteer-only model.

## Appendix F: Validation pack v1

---

Source file: `validation_pack_v1.md`

# MCS Validation Pack v1

---

**Companion artefact** to *A Supervisory-Evidence Ontology for Agentic AI under EU Law: Minimum Conceptual Set and Temporal Extension* (Janssen, April 2026).

**Filename:** `validation_pack_v1.md` **Purpose:** pre-registered open-call mechanics for three independent volunteer validation tracks operating against the MCS specification artefacts. **Licence:** CC BY 4.0 **Status:** open call. All three tracks accept volunteer submissions from publication date.

---

## §1. Why three tracks

---

The paper's central design claim is that the MCS synthesis is intended to be *defensible, reproducible, and testably better than ad-hoc practice*. Each adjective is a separate empirical hypothesis and requires a separate test. This pack is a validation design, not completed validation evidence.

- **Defensible** is tested by inter-rater consistency on the empirical case sample. If two independent coders applying the same rubric reach  $\kappa < 0.60$ , the claim that the synthesis is intersubjectively coherent fails. This is **Track 1**.
- **Materially better** is tested by technical feasibility at production-relevant throughput. If SHACL validation latency exceeds operational thresholds at realistic decision rates, the operational-usability hypothesis for MCS fails. This is **Track 2**.
- **Reproducible across topologies** is tested by structural fit against agent topologies the working paper has not stress-tested. If a contributor with a different topology cannot represent their event stream within the 23-class MCS, the structural-completeness claim fails for that topology and the gap becomes a research item. This is **Track 3**.

Failure on any of the three is a publishable result. The validation pack does not pre-suppose success; it pre-registers the tests.

---

## §2. Volunteer model

---

This validation programme operates on volunteer participation. There is no compensation, no honorarium, no advisory time exchange, and no other consideration offered to volunteers beyond co-citation in subsequent paper revisions where applicable.

The volunteer model is deliberate. Compensation arrangements would either compromise coder independence (in Track 1) or signal commercial relationships (in Tracks 2 and 3) that the open-call mechanic is designed to avoid. Volunteers participate because the topic interests them, because they want to contribute to a defensible governance specification, or because their organisation benefits from validating its own evidentiary infrastructure against an external reference.

The volunteer keeps full ownership of any data they submit. The author has no claim on volunteer measurement data, coder rationale notes, or topology documentation beyond the right to cite findings under CC BY 4.0 in subsequent paper revisions.

---

## §3. Track 1: Independent replication of case classification

---

### Hypothesis (pre-registered)

A volunteer applying the \$6 rubric of `coding_rubrics_v0_2.md` to the 25-case sample in `mcs_case_sample_v0_2.csv` reaches Cohen's  $\kappa \geq 0.70$  against the round-1 codings on primary classification.

Failure criterion:  $\kappa < 0.60$  triggers rubric revision rather than cell-level reconciliation. The failure itself is part of the publication commitment.

### Materials provided

1. `coding_rubrics_v0_2.md` — full reproducibility protocols.
2. `mcs_case_sample_v0_2.csv` with `primary_classification`, `finding1_bucket`, `sensitivity_flag`, `grounds_NonQualifyingADM_subclass`, and `supports_finding` columns redacted.
3. The 25 case records (jurisdiction, date, brief facts, source citation) in unredacted form.

### Volunteer requirements

- LL.M. or equivalent in EU data protection law, or equivalent professional experience.
- Working familiarity with WP251rev.01, SCHUFA, Dun and Bradstreet Austria, AP handvatten July 2025.
- No prior involvement with the paper's authorship or review.

### Procedure

1. Volunteer requests blinded materials by emailing the contact address in §9 with a conflict-of-interest disclosure and a brief statement of qualifications.
2. Volunteer applies the five-step protocol (`coding_rubrics_v0_2.md` §6) to each of 25 cases, in random order with seed disclosed at submission.

3. Volunteer records per-case Step 1-5 outcomes plus primary classification plus sensitivity flag plus NonQualifyingADM subclass mapping in the format specified by the CSV schema.
4. Volunteer submits the completed CSV plus a short coder-disclosure document covering background, conflicts of interest, and methodological notes.

There is no fixed deadline. Volunteers proceed at their own pace. The author commits to publishing the  $\kappa$  result in the next paper revision after the first qualifying submission arrives.

### Computation

For each artefact column, Cohen's  $\kappa$  is computed against the round-1 codings. Per-cell disagreements are catalogued with both codings retained.

Reported metrics: -  $\kappa$  on primary classification. -  $\kappa$  on `finding1_bucket`. -  $\kappa$  on `sensitivity_flag`. -  $\kappa$  on NonQualifyingADM subclass mapping (where applicable). - Per-Finding sensitivity ranges under disagreement-flipping.

### Failure handling

- $\kappa < 0.60$ : rubric revision plus re-coding by a third independent volunteer against revised rubrics.
- $\kappa \geq 0.60$  but  $< 0.70$ : per-artefact  $\kappa$  reported; cells with disagreement noted in publication metadata.
- $\kappa \geq 0.70$ : substantial agreement reported; original codings retained as canonical.

### Recognition

The volunteer is credited as Track 1 coder in the next paper revision and any subsequent peer-reviewed publication. The volunteer retains full ownership of the rationale notes and may publish their own analysis independently.

---

## §4. Track 2: Technical feasibility benchmark

---

### Hypothesis (pre-registered)

A volunteer emitting MCS-conformant Turtle at decision rates of 1, 10, 100, and 1000 decisions per second achieves p95 SHACL validation latency below 100 ms at 100 decisions per second on at least one of three reference SHACL engines: `pyshacl`, Apache Jena SHACL, TopBraid SHACL API.

Failure criterion: p95 latency exceeds 100 ms at 100 decisions per second on all three engines, indicating that operational deployment requires either MCS shape simplification, sharding, or non-SHACL validation paths.

### Materials provided

1. `mcs_vocabulary.ttl` — the 23-class OWL/RDF vocabulary.
2. `mcs_profiles_v1_2.ttl` — Profile A plus Profile B SHACL shapes.
3. The 11 core temporal-extension shapes referenced in §6.3 of the paper.

### Volunteer requirements

- Working AI agent system or equivalent test harness capable of generating decision events at the specified rates.
- Technical capacity to populate Turtle and run SHACL validation.
- Documentation of the agent topology including decision granularity, event stream characteristics, and runtime infrastructure.

### Procedure

1. Volunteer populates MCS-conformant Turtle from their event stream at the specified decision rates.
2. Volunteer runs SHACL validation against `mcs_profiles_v1_2.ttl` using each of the three reference engines.
3. Volunteer measures p50, p95, and p99 latency at each rate-engine combination.
4. Volunteer submits:
  - Latency measurements per rate per engine.
  - Topology documentation.
  - SHACL engine versions used.
  - Hardware specifications (CPU, RAM, storage).
  - The Turtle workload sample (or a representative subset).

There is no fixed deadline. Volunteers proceed at their own pace.

### Computation

For each rate-engine combination: latency distribution reported with confidence intervals. Aggregation across volunteers compares latency variance attributable to topology versus engine versus hardware.

Reported metrics: - Latency per (rate, engine) combination. - Topology dependency analysis. - Failure modes documented (timeouts, memory exhaustion, validation errors).

## Failure handling

- p95 latency below 100 ms at 100 decisions per second: hypothesis confirmed; technical feasibility established for that volunteer's topology.
- p95 latency between 100 ms and 1 second: partial confirmation; specification recommendations issued for shape simplification.
- p95 latency above 1 second: hypothesis fails for that volunteer; gap becomes research item for shape redesign or non-SHACL validation paths.

## Recognition

Volunteers are credited as Track 2 contributors in the next paper revision and any subsequent peer-reviewed publication. Volunteers retain full ownership of their measurement data and may publish independently.

---

## §5. Track 3: Structural fit pilot

---

### Hypothesis (pre-registered)

A volunteer with an agent topology distinct from those represented in the paper's seven worked scenarios ( `mcs_scenarios_v0_2.md` ) can represent every event class in their event stream using an MCS class without requiring extension.

Failure criterion: at least one event class requires a new MCS class that is not subclass-of an existing slot. The failure is itself the research finding: the gap is documented in the next paper revision as a structural-completeness limitation.

### Materials provided

1. `mcs_vocabulary.ttl` — 23-class declaration.
2. `mcs_framework_specification.md` — human-readable equivalent.
3. The seven worked scenarios in `mcs_scenarios_v0_2.md` for reference.

### Volunteer requirements

- Agent topology not represented in the seven worked scenarios. Examples of distinct topologies that would qualify:
  - Multi-agent fleet with peer-to-peer messaging.
  - Insurance-claims agent with sectoral overlay (e.g., MiFID II or MDR Rule 11).
  - Cross-jurisdictional agent operating under simultaneous AP-permissive and CNIL-strict supervisory regimes.
  - Agent operating in a Member State with non-EU regulatory overlay (e.g., UK ICO post-Brexit, Italy Law 132/2025).
- Documented event-class taxonomy of the topology.
- Capacity to run a structural mapping exercise from event class to MCS class.

### Procedure

1. Volunteer enumerates their event-class taxonomy.
2. Volunteer maps each event class to an MCS class (Slots 1-23).
3. Volunteer identifies any event class that does not fit, with rationale.
4. Volunteer submits:
  - Topology description.
  - Event-class taxonomy.
  - Event-class to MCS-class mapping.
  - Gap report (if any).
  - Recommendations for shape extension or new slot creation.

There is no fixed deadline. Volunteers proceed at their own pace.

### Computation

Gap reports are aggregated across volunteers. Recurring gaps become candidate slots for vocabulary updates; isolated gaps become topology-specific extensions.

Reported metrics: - Number of event classes mapped per volunteer. - Number of gaps reported per volunteer. - Common gap patterns across volunteers. - Recommended structural extensions.

### Failure handling

- Zero gaps reported: structural sufficiency confirmed for that volunteer's topology.
- Gaps reported but resolvable by subclass-of relationships within existing slots: minor structural extension proposed.
- Gaps requiring new slots: documented as a research item with the volunteer credited as co-author of the relevant section in subsequent revisions.

### Recognition

Same as Track 2: co-citation in subsequent paper revisions, no financial compensation, volunteer retains data ownership.

---

## §6. Open submission and eligibility

---

Any volunteer meeting the per-track requirements may submit. Submissions are evaluated on a rolling basis. The author commits to acknowledging receipt and providing a substantive response.

### Required submission documents

For all three tracks: 1. Volunteer identity and affiliation (or pseudonym with verifiable credentials). 2. Conflict-of-interest disclosure. 3. Track-specific deliverables as specified in §3-§5. 4. Licensing declaration: the volunteer agrees that submission contents may be cited under CC BY 4.0 in subsequent paper revisions and any related publication.

### Anonymisation option

Volunteers operating in regulated sectors where public attribution would expose competitive or compliance positions may submit anonymously. Anonymous submissions are accepted on the condition that the author has visibility on volunteer identity for verification purposes; published metadata redacts identifying information at the volunteer's request.

### Disqualification grounds

- Failure to disclose conflicts of interest.
- Submission of fabricated data.
- Violation of the per-track procedure (e.g., Track 1 coder receiving unredacted classifications prior to coding).

---

## §7. Aggregation and publication

---

### Subsequent paper revisions

The next paper revision will incorporate Track 1 results as a new §4.6.5 (Inter-rater reliability report) once a qualifying Track 1 submission arrives. Tracks 2 and 3 results will be incorporated as available; if no Track 2 or Track 3 submissions arrive before a revision, the section will note the absence rather than delay publication.

Track 2 and Track 3 results arriving after a revision will be incorporated into vocabulary updates. The MCS specification artefacts (`mcs_vocabulary.ttl`, `mcs_profiles_v1_2.ttl`) may be updated based on volunteer findings; updates are versioned and changelogs.

### Aggregation procedure

For each track: 1. Submissions catalogued in a contributor registry maintained at the project repository. 2. Per-submission metrics reported. 3. Cross-submission aggregation reported where contributor count permits (target: at least three submissions per track for cross-aggregation). 4. Publication of aggregate findings under CC BY 4.0.

### Negative results

Negative results (κ failures, latency failures, structural-fit failures) are published with the same prominence as positive results. The validation pack pre-commits to publishing failure reports as integral parts of subsequent revisions rather than as supplementary footnotes.

---

## §8. Limits of the validation pack

---

### What the pack produces

- Independent reliability evidence on the case-classification artefact (Track 1).
- Technical feasibility measurements on production-relevant SHACL throughput (Track 2).
- Structural fit findings across agent topologies not represented in the paper (Track 3).

### What the pack does not produce

- **Regulatory validation.** No DPA has endorsed the MCS; no CJEU judgment has applied it; no notified body under the AI Act has cross-mapped it. Regulatory validation is a separate programme requiring engagement with EDPB, the European AI Office, CEN-CENELEC JTC 21, and W3C DPVCG. The validation pack does not substitute for regulatory engagement.
- **Population-level claims.** The 25-case empirical sample is selection-biased toward litigated cases. Track 1 cannot correct for selection bias; it can only test inter-rater consistency on the existing biased sample.
- **Causal claims about supervisory outcomes.** The pack does not test whether MCS-structured evidence changes regulatory determinations relative to ad-hoc evidence. That requires a controlled comparison study with supervisor participation, which is outside the validation pack's scope.

### What the pack avoids

- Cell-level reconciliation between coders. The pack pre-commits to rubric revision on  $\kappa$  failure rather than horse-trading individual disagreements into apparent consensus.
- Vendor or commercial endorsement. The pack is open to any qualifying volunteer; no volunteer is preferred on commercial grounds.
- Author-side intervention in coder reasoning. Track 1 coders work blind; the author does not engage with coding decisions during the coding phase.
- Compensation in any form. Volunteer-only participation eliminates the perverse-incentive problem at source.

### **Self-selection bias of the volunteer pool**

Volunteer-only participation produces self-selection bias. The volunteer pool is likely to overrepresent ontology-literate, governance-friendly, or author-sympathetic contributors, and to underrepresent practitioners who are skeptical of formal-ontology approaches, time-constrained operational legal counsel, and contributors from jurisdictions where the SCHUFA path is contested or rejected. The bias is acknowledged here rather than corrected, because the alternative (paid recruitment of a representative panel) introduces a different and worse incentive distortion. Mitigation is open submission across all three tracks plus the pre-commitment to publishing negative results with the same prominence as positive ones (see §7). Readers of any aggregate  $\kappa$ , latency, or structural-fit number produced by this pack should treat it as evidence from a self-selected pool, not as a population estimate.

---

### **§9. Repository and contact**

The MCS specification artefacts plus this validation pack are maintained at the project repository (Zenodo DOI to be assigned at publication; arXiv preprint identifier to be assigned at submission). Contact for submissions and queries: through the contact channel disclosed at [apparens.nl](https://apparens.nl).

---

*End of `validation_pack_v1.md`. © 2026 J. Janssen, Apparens. CC BY 4.0.*

# Appendix G: Vocabulary (Turtle)

Source file: `mcs_vocabulary.ttl`

```
# =====
# MCS Vocabulary – Minimum Conceptual Set for Supervisory Evidence
# =====
#
# Companion artefact to:
# "A Supervisory-Evidence Ontology for Agentic AI under EU Law:
# Minimum Conceptual Set and Temporal Extension"
# Jeroen Janssen, Apparens, April 2026
#
# This file declares the 23 MCS classes specified in Section 4.3 of the paper,
# split architecturally into:
# - mcs-core (Table 4.3a): twelve classes covering agent action primitives,
# decision and intervention semantics, and delegation-and-policy temporal
# structure (slots 1-12).
# - mcs-supervisory (Table 4.3b): eleven classes covering sectoral regime
# pointers, supervisory scope, typed log events, monitoring observations,
# drift indicators, evidence artefacts, serious incidents, and the
# NonQualifyingADM residual (slots 13-23).
#
# Five DPV-imported classes (DataSubject, PersonalData, Purpose, LegalBasis,
# Consent) are reused via owl:imports rather than redeclared.
#
# A planned split into mcs_core.ttl and mcs_supervisory.ttl will operationalise
# the architectural layering as distinct owl:imports targets. This file is
# the consolidated declaration pending that split.
#
# -----
# Disclaimer (paper §6.3 and §7.3)
# -----
# Structural-completeness vocabulary only; not a legal-compliance instrument. Validation against
# shapes derived
# from this vocabulary (e.g. mcs_profiles_v1_2.ttl) yields sh:conforms=true
# or sh:conforms=false reflecting structural conformance with the named
# profile's shapes. It does NOT establish compliance with the AI Act, the
# GDPR, NIS2, DORA, the CRA, or any other instrument of Union or national
# law. Structural completeness is necessary but not sufficient for
# supervisory documentation obligations. A supervisor who reads only the
# SHACL conforms flag without reading the underlying evidence is not
# substituting MCS validation for legal judgement; they are misusing the
# tool. The disclaimer is not self-enforcing.
#
# -----
# Licensing and registration
# -----
# © 2026 J. Janssen, Apparens
# Licence: CC BY 4.0
# i-DEPOT: 158508 (AI Enterprise Control Index)
# Paper reference: ISBN 9798252058993 (English) and Dutch edition
#
# =====

@prefix mcs:      <http://apparens.nl/mcs#> .
@prefix gov:      <http://apparens.nl/mcs/gov#> .
@prefix owl:    <http://www.w3.org/2002/07/owl#> .
@prefix rdf:      <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix rdfs:     <http://www.w3.org/2000/01/rdf-schema#> .
@prefix xsd:      <http://www.w3.org/2001/XMLSchema#> .
@prefix skos:     <http://www.w3.org/2004/02/skos/core#> .
@prefix dcterms: <http://purl.org/dc/terms/> .
@prefix prov:     <http://www.w3.org/ns/prov#> .
@prefix time:     <http://www.w3.org/2006/time#> .
@prefix dpv:      <https://w3id.org/dpv#> .
@prefix gdprov:   <https://opscience.adaptcentre.ie/ontologies/GDPProv#> .
@prefix eu-aiact: <https://w3id.org/dpv/legal/eu/aiact#> .

# =====
# Ontology header
# =====

<http://apparens.nl/mcs> a owl:Ontology ;
  dcterms:title "Minimum Conceptual Set for Supervisory Evidence (MCS)"@en ;
  dcterms:creator "Jeroen Janssen, Apparens" ;
  dcterms:date "2026-04-25"^^xsd:date ;
  dcterms:license <https://creativecommons.org/licenses/by/4.0/> ;
  dcterms:rights "© 2026 J. Janssen Apparens, CC BY 4.0; i-DEPOT 158508" ;
  rdfs:comment ""The MCS specifies 23 conceptual slots / net-new classes under the paper's counting
  convention (12 core + 11 supervisory)
  anchored in operative provisions of GDPR, AI Act, and NIS2, plus
  structured motivation against a 25-case EU ADM enforcement sample.
  Structural-completeness vocabulary; not a legal-compliance instrument.""@en ;
  owl:imports <http://www.w3.org/ns/prov-o> ,
    <http://www.w3.org/2006/time> ,
    <https://w3id.org/dpv> .
  # gdprov and eu-aiact imports are SOFT (referenced via prefix where
  # available) because their canonical IRIs are not consistently
  # resolvable as of April 2026.

# =====
# mcs-core: agent-behaviour layer (Table 4.3a, slots 1-12)
# =====

# -----
# Slot 1: gov:Agent (with subtypes SoftwareAgent, HumanAgent, OrganizationalAgent)
# Anchor: GDPR Art. 4(7)-(8); AI Act Art. 3(3)-(4); SCHUFA para 48
# -----
```

```

gov:Agent a owl:Class ;
  rdfs:subClassOf prov:Agent ;
  rdfs:label "Supervisory-scope Agent"@en ;
  rdfs:comment """"Abstract supervisory-scope agent. Covers software agents
    (LLM-based or symbolic), human agents acting in operational capacity,
    and organisational agents (controllers, processors, deployers).
    Subclass of prov:Agent to inherit attribution semantics.""@en ;
  skos:note "Paper §4.3 slot 1; Table 4.3a." .

gov:SoftwareAgent a owl:Class ;
  rdfs:subClassOf gov:Agent , prov:SoftwareAgent ;
  rdfs:label "Software Agent"@en ;
  rdfs:comment "An autonomous or semi-autonomous software process, including LLM-based agents." .

gov:HumanAgent a owl:Class ;
  rdfs:subClassOf gov:Agent , prov:Person ;
  rdfs:label "Human Agent"@en ;
  rdfs:comment "A natural person acting in operational capacity within a controller or processor." .

gov:OrganizationalAgent a owl:Class ;
  rdfs:subClassOf gov:Agent , prov:Organization ;
  rdfs:label "Organisational Agent"@en ;
  rdfs:comment "A legal entity acting as controller, processor, deployer, or other AI Act/GDPR role-
  bearer." .

# -----
# Slot 2: gov:LegalRole (eight AI Act and GDPR subtypes)
# Anchor: GDPR Arts 4(7)-(8), 26, 28; AI Act Art. 3(3)-(8)
# -----

gov:LegalRole a owl:Class ;
  rdfs:label "Legal Role"@en ;
  rdfs:comment """"Abstract legal role under GDPR or AI Act. Time-indexed
  attribution to gov:Agent via gov:RoleAssignment. A single
  organisation can simultaneously hold dpv:Controller under GDPR and
  gov:Deployer under the AI Act; the assignments are distinct.""@en ;
  skos:note "Paper §4.3 slot 2; §4.2 Legal Role Triad." .

gov:Provider a owl:Class ;
  rdfs:subClassOf gov:LegalRole ;
  rdfs:label "Provider (AI Act)"@en ;
  rdfs:comment "AI Act Art. 3(3): natural or legal person developing an AI system." ;
  rdfs:seeAlso eu-aiact:Provider .

gov:Deployer a owl:Class ;
  rdfs:subClassOf gov:LegalRole ;
  rdfs:label "Deployer (AI Act)"@en ;
  rdfs:comment "AI Act Art. 3(4): natural or legal person using an AI system." ;
  rdfs:seeAlso eu-aiact:Deployer .

gov:Distributor a owl:Class ;
  rdfs:subClassOf gov:LegalRole ;
  rdfs:label "Distributor (AI Act)"@en ;
  rdfs:comment "AI Act Art. 3(7)." .

gov:Importer a owl:Class ;
  rdfs:subClassOf gov:LegalRole ;
  rdfs:label "Importer (AI Act)"@en ;
  rdfs:comment "AI Act Art. 3(6)." .

gov:Manufacturer a owl:Class ;
  rdfs:subClassOf gov:LegalRole ;
  rdfs:label "Product Manufacturer (AI Act)"@en ;
  rdfs:comment "AI Act Art. 3(8); product manufacturer integrating AI." .

gov:AuthorisedRepresentative a owl:Class ;
  rdfs:subClassOf gov:LegalRole ;
  rdfs:label "Authorised Representative (AI Act)"@en ;
  rdfs:comment "AI Act Art. 3(5)." .

gov:Controller a owl:Class ;
  rdfs:subClassOf gov:LegalRole ;
  rdfs:label "Controller (GDPR)"@en ;
  rdfs:comment "GDPR Art. 4(7)." ;
  rdfs:seeAlso dpv:Controller .

gov:Processor a owl:Class ;
  rdfs:subClassOf gov:LegalRole ;
  rdfs:label "Processor (GDPR)"@en ;
  rdfs:comment "GDPR Art. 4(8); subject to Art. 28 contract requirements." ;
  rdfs:seeAlso dpv:Processor .

gov:RoleAssignment a owl:Class ;
  rdfs:subClassOf prov:Attribution ;
  rdfs:label "Role Assignment"@en ;
  rdfs:comment """"Time-indexed attribution of a gov:LegalRole to a gov:Agent.
  Allows reconstruction of role ownership at any past moment. Required
  properties: gov:hasAgent, gov:hasRole, gov:assignmentInterval.""@en .

gov:hasAgent a owl:ObjectProperty ;
  rdfs:domain gov:RoleAssignment ;
  rdfs:range gov:Agent ;
  rdfs:label "has agent"@en .

gov:hasRole a owl:ObjectProperty ;
  rdfs:domain gov:RoleAssignment ;
  rdfs:range gov:LegalRole ;
  rdfs:label "has role"@en .

gov:assignmentInterval a owl:ObjectProperty ;
  rdfs:domain gov:RoleAssignment ;
  rdfs:range time:ProperInterval ;
  rdfs:label "assignment interval"@en .

```

```

# -----
# Slot 3: gov:Activity (agent action; PROV-O reuse)
# Anchor: AI Act Art. 12
# -----

gov:Activity a owl:Class ;
  rdfs:subClassOf prov:Activity ;
  rdfs:label "Agent Activity"@en ;
  rdfs:comment ""Agent action within MCS supervisory scope. Inherits
    prov:startedAtTime, prov:endedAtTime, prov:wasAssociatedWith.
    Counted as zero net new classes under §2.5 counting convention
    (direct PROV-O reuse with no new properties).""@en ;
  skos:note "Paper §4.3 slot 3." .

# -----
# Slot 4: gov:ToolInvocation
# Anchor: AI Act Art. 13
# -----

gov:ToolInvocation a owl:Class ;
  rdfs:subClassOf gov:Activity ;
  rdfs:label "Tool Invocation"@en ;
  rdfs:comment "A specific tool-call activity performed by an agent. Anchored in AI Act Art. 13
    transparency obligations." ;
  skos:note "Paper §4.3 slot 4; §5.3 sketch F3." .

gov:hasTool a owl:ObjectProperty ;
  rdfs:domain gov:ToolInvocation ;
  rdfs:label "has tool"@en ;
  rdfs:comment "Reference to the tool, function, or external service invoked." .

gov:hasInvocationInput a owl:ObjectProperty ;
  rdfs:domain gov:ToolInvocation ;
  rdfs:label "has invocation input"@en .

gov:hasInvocationOutput a owl:ObjectProperty ;
  rdfs:domain gov:ToolInvocation ;
  rdfs:label "has invocation output"@en .

# -----
# Slot 5: gov:MemoryEvent (with MemoryRead, MemoryWrite)
# Anchor: AI Act Art. 12; Art. 72
# -----

gov:MemoryEvent a owl:Class ;
  rdfs:subClassOf gov:Activity ;
  rdfs:label "Memory Event"@en ;
  rdfs:comment "Provenance of model or agent memory access; abstract over read and write." ;
  skos:note "Paper §4.3 slot 5." .

gov:MemoryRead a owl:Class ;
  rdfs:subClassOf gov:MemoryEvent ;
  rdfs:label "Memory Read"@en ;
  rdfs:comment "An agent reads from memory. Source for §5.3 sketch F5 (research conjecture on Art.
    22 contributing factor)." .

gov:MemoryWrite a owl:Class ;
  rdfs:subClassOf gov:MemoryEvent ;
  rdfs:label "Memory Write"@en ;
  rdfs:comment "An agent writes to memory. Subject to NIS2 Art. 23(4) preservation requirement
    during incident windows; see App A.15 sketch F12." .

gov:hasMemoryReference a owl:DatatypeProperty ;
  rdfs:domain gov:MemoryEvent ;
  rdfs:range xsd:anyURI ;
  rdfs:label "memory reference"@en ;
  rdfs:comment "URI or addressable identifier for the memory location." .

# -----
# Slots 6-10: DPV reuse (not redeclared)
# Slot 6: dpv:DataSubject - GDPR Art. 4(1), 15, 22
# Slot 7: dpv:PersonalData - GDPR Art. 9, 10, 30(1)(c)
# Slot 8: dpv:Purpose - GDPR Art. 5(1)(b), 6(4), 30
# Slot 9: dpv:LegalBasis - GDPR Art. 6, 9
# Slot 10: dpv:Consent + gdprov consent lifecycle - GDPR Art. 6(1)(a), 7
# -----
# Per Section 2.5 counting convention, these five are zero net new MCS
# classes (direct DPV reuse). The MCS uses them via owl:imports above.
# -----
# Slot 11: gov:Decision (three subtypes plus gov:Score plus drawsStronglyOn)
# Anchor: GDPR Art. 22(1); SCHUFA para 73; Dun and Bradstreet para 40
# -----

gov:Decision a owl:Class ;
  rdfs:label "Decision (Article 22 taxonomy)"@en ;
  rdfs:comment ""Decision concerning a data subject. Three-way taxonomy
    anchored in GDPR Art. 22(1) and SCHUFA upstream-attribution path:
    SolelyAutomatedDecision, AssistedDecisionWithMeaningfulHumanIntervention,
    NonAutomatedDecision. Profile-dependent classification; see paper §4.6
    Findings 3 and 4.""@en ;
  skos:note "Paper §4.3 slot 11; §5.3 sketch F1; §6.4 dual operationalisation." .

gov:SolelyAutomatedDecision a owl:Class ;
  rdfs:subClassOf gov:Decision ;
  rdfs:label "Solely Automated Decision"@en ;
  rdfs:comment "Decision based solely on automated processing within Art. 22(1) scope. Profile-
    dependent boundary." .

gov:AssistedDecisionWithMeaningfulHumanIntervention a owl:Class ;
  rdfs:subClassOf gov:Decision ;
  rdfs:label "Assisted Decision (with Meaningful Human Intervention)"@en ;

```

```

    rdfs:comment "Decision with human intervention meeting WP251rev.01 plus AP handvatten conjunctive
criteria." .

gov:NonAutomatedDecision a owl:Class ;
    rdfs:subClassOf gov:Decision ;
    rdfs:label "Non-Automated Decision"@en ;
    rdfs:comment "Decision without ADM substrate; outside Art. 22 scope by absence of automated
processing." .

gov:Score a owl:Class ;
    rdfs:label "Score"@en ;
    rdfs:comment """"Preparatory automated probability or risk score. Per SCHUFA
para 73, a score is itself a decision under Art. 22(1) where downstream
actors draw strongly on it. See paper §5.3 sketch F2.""@en .

gov:drawsStronglyOn a owl:ObjectProperty ;
    rdfs:domain gov:Decision ;
    rdfs:range gov:Score ;
    rdfs:label "draws strongly on"@en ;
    rdfs:comment """"Profile-dependent predicate (paper §6.4). Profile A
(AP-permissive, quantitative): flip-rate at  $\epsilon=0.10$  exceeds  $M=0.50$ .
Profile B (CNIL-strict, qualitative): any of three judicial indicia
(specialised scoring model; refusal frequency  $\geq 0.90$ ; absent
documented intervention).""@en .

gov:hasLegalEffect a owl:DatatypeProperty ;
    rdfs:domain gov:Decision ;
    rdfs:range xsd:boolean ;
    rdfs:label "has legal or significant effect"@en ;
    rdfs:comment "Art. 22(1) qualifier: produces legal effects concerning the data subject or
similarly significantly affects them." .

gov:flipRateAtEpsilon10 a owl:DatatypeProperty ;
    rdfs:domain gov:Score ;
    rdfs:range xsd:decimal ;
    rdfs:label "flip rate at  $\epsilon=0.10$ "@en ;
    rdfs:comment "Profile A quantitative property. Empirical proportion of decisions that flip outcome
under 10% perturbation of score." .

gov:usesSpecialisedScoringModel a owl:DatatypeProperty ;
    rdfs:domain gov:Score ;
    rdfs:range xsd:boolean ;
    rdfs:label "uses specialised scoring model"@en ;
    rdfs:comment "Profile B qualitative indicium 1. Wiesbaden VG January 2026: existence of
specialised scoring model is itself evidence of strong reliance." .

gov:refusalFrequencyGivenScore a owl:DatatypeProperty ;
    rdfs:domain gov:Decision ;
    rdfs:range xsd:decimal ;
    rdfs:label "refusal frequency given insufficient score"@en ;
    rdfs:comment "Profile B qualitative indicium 2. Operationalises SCHUFA para 48 'in almost all
cases' factual finding." .

gov:hasDocumentedIntervention a owl:DatatypeProperty ;
    rdfs:domain gov:Decision ;
    rdfs:range xsd:boolean ;
    rdfs:label "has documented human intervention"@en ;
    rdfs:comment "Profile B qualitative indicium 3. Operationalises D&B para 40 'lacked any manual
oversight'." .

# -----
# Slot 12: gov:HumanIntervention (six-property meaningful-intervention schema)
# Anchor: GDPR Art. 22(3); AI Act Art. 14; WP251rev.01; AP handvatten July 2025
# -----

gov:HumanIntervention a owl:Class ;
    rdfs:label "Human Intervention"@en ;
    rdfs:comment """"Six-property schema operationalising the WP251rev.01 plus
AP handvatten meaningful-intervention test, with the
authorityActuallyExercised property added per Uber/Ola Hof Amsterdam
(B10, ECLI:NL:GHAMS:2023:793). All six properties must hold for an
intervention to count as meaningful under the conjunctive test.""@en ;
    skos:note "Paper §4.3 slot 12; §4.6 Finding 3; App A.12 sketch F9." .

gov:hasAuthorityToDeviate a owl:DatatypeProperty ;
    rdfs:domain gov:HumanIntervention ;
    rdfs:range xsd:boolean ;
    rdfs:label "has authority to deviate"@en ;
    rdfs:comment "Property 1: reviewer holds documented authority to override the automated outcome."
.

gov:hasUnderstandingOfLogic a owl:DatatypeProperty ;
    rdfs:domain gov:HumanIntervention ;
    rdfs:range xsd:boolean ;
    rdfs:label "has understanding of logic"@en ;
    rdfs:comment "Property 2: reviewer holds documented understanding of the system's decision logic
per Art. 15(1)(h)." .

gov:consideredAllRelevantData a owl:DatatypeProperty ;
    rdfs:domain gov:HumanIntervention ;
    rdfs:range xsd:boolean ;
    rdfs:label "considered all relevant data"@en ;
    rdfs:comment "Property 3: reviewer documented consideration of relevant case-specific factors." .

gov:hasMarginOfDiscretion a owl:DatatypeProperty ;
    rdfs:domain gov:HumanIntervention ;
    rdfs:range xsd:boolean ;
    rdfs:label "has margin of discretion"@en ;
    rdfs:comment "Property 4: reviewer's role permits substantive discretion, not merely rubber-
stamping." .

gov:interventionDurationSeconds a owl:DatatypeProperty ;
    rdfs:domain gov:HumanIntervention ;

```

```

rdfs:range xsd:nonNegativeInteger ;
rdfs:label "intervention duration (seconds)"@en ;
rdfs:comment ""Property 5: time spent on intervention. Default operational
threshold ≥60s for non-symbolic review, sector-calibrated (5min for
complex clinical triage; 15s for high-volume fraud-detection alerts
conditional on case-specific examination). See paper §2.7.""@en .

gov:authorityActuallyExercised a owl:DatatypeProperty ;
rdfs:domain gov:HumanIntervention ;
rdfs:range xsd:boolean ;
rdfs:label "authority actually exercised"@en ;
rdfs:comment ""Property 6: per Uber/Ola Hof Amsterdam (B10), authority to
deviate is necessary but not sufficient. The authority must be
actually exercised in practice; population-level deviation rate
> 0 (illustrative 0=10%). Profile B requires this property=true
for MeaningfulIntervention.""@en .

gov:contributesTo a owl:ObjectProperty ;
rdfs:domain gov:HumanIntervention ;
rdfs:range gov:Decision ;
rdfs:label "contributes to"@en .

# =====
# mcs-supervisory: supervisory-evidence layer (Table 4.3b, slots 13-23)
# =====

# -----
# Slot 13: gov:Policy, gov:PolicyVersion
# Anchor: AI Act Art. 9 risk management; Art. 17 QMS
# -----

gov:Policy a owl:Class ;
rdfs:label "Policy"@en ;
rdfs:comment "An organisational or technical policy governing agent behaviour." ;
skos:note "Paper §4.3 slot 13." .

gov:PolicyVersion a owl:Class ;
rdfs:subClassOf prov:Entity ;
rdfs:label "Policy Version"@en ;
rdfs:comment ""Versioned snapshot of a Policy with temporal validity and
supersession. PolicyVersionShape (mcs_profiles_v1_2.ttl) requires
non-overlapping validity intervals across successor versions.""@en .

gov:supersedesPolicy a owl:ObjectProperty ;
rdfs:domain gov:PolicyVersion ;
rdfs:range gov:PolicyVersion ;
rdfs:label "supersedes policy"@en .

gov:isMorePermissive a owl:DatatypeProperty ;
rdfs:domain gov:PolicyVersion ;
rdfs:range xsd:boolean ;
rdfs:label "is more permissive"@en ;
rdfs:comment "Used in App A.10 sketch F7 (lex mitior conjecture, RC-flagged)." .

gov:policyValidityInterval a owl:ObjectProperty ;
rdfs:domain gov:PolicyVersion ;
rdfs:range time:ProperInterval ;
rdfs:label "policy validity interval"@en .

# -----
# Slot 14: gov:RiskClassification (four AI Act subtypes)
# Anchor: AI Act Art. 5, 6, 50, Annex III; ISO/IEC 23894
# -----

gov:RiskClassification a owl:Class ;
rdfs:label "Risk Classification"@en ;
rdfs:comment "AI Act risk classification of an AI system." ;
skos:note "Paper §4.3 slot 14; App A.13 sketch F10." .

gov:ProhibitedAI a owl:Class ;
rdfs:subClassOf gov:RiskClassification ;
rdfs:label "Prohibited AI"@en ;
rdfs:comment "AI Act Art. 5 prohibited practice." .

gov:HighRiskAI a owl:Class ;
rdfs:subClassOf gov:RiskClassification ;
rdfs:label "High-Risk AI"@en ;
rdfs:comment "AI Act Art. 6 plus Annex III high-risk system. Subject to Chapter III obligations."
.

gov:LimitedRiskAI a owl:Class ;
rdfs:subClassOf gov:RiskClassification ;
rdfs:label "Limited-Risk AI"@en ;
rdfs:comment "AI Act Art. 50 transparency obligations only." .

gov:MinimalRiskAI a owl:Class ;
rdfs:subClassOf gov:RiskClassification ;
rdfs:label "Minimal-Risk AI"@en ;
rdfs:comment "Outside the scope of Arts. 5, 6, and 50; voluntary codes of conduct under Art. 95."
.

# -----
# Slot 15: gov:IntendedPurpose
# Anchor: AI Act Art. 3(12), 13
# -----

gov:IntendedPurpose a owl:Class ;
rdfs:label "Intended Purpose (AI Act)"@en ;
rdfs:comment ""AI Act Art. 3(12): use for which the AI system is intended
by the provider, including specific context and conditions of use, as
specified in instructions for use. Distinct from dpv:Purpose (which is
the GDPR Art. 5(1)(b) purpose-of-processing concept).""@en ;
skos:note "Paper §4.3 slot 15." .

```

```

# -----
# Slot 16: gov:SectoralRegime (with named instances)
# Anchor: Sectoral statutes
# -----

gov:SectoralRegime a owl:Class ;
  rdfs:label "Sectoral Regime"@en ;
  rdfs:comment "Pointer to a sectoral regulatory regime applicable to the AI system." ;
  skos:note "Paper §4.3 slot 16; §7.6 regulatory-coherence assumption." .

gov:competentSupervisor a owl:ObjectProperty ;
  rdfs:domain gov:SectoralRegime ;
  rdfs:range gov:SupervisoryScope ;
  rdfs:label "competent supervisor"@en .

# Named instances of sectoral regimes
gov:MiFID_II a gov:SectoralRegime ;
  rdfs:label "MiFID II"@en ;
  rdfs:comment "Directive 2014/65/EU on markets in financial instruments." .

gov:MDR a gov:SectoralRegime ;
  rdfs:label "MDR"@en ;
  rdfs:comment "Regulation (EU) 2017/745 on medical devices. Rule 11 governs autonomous clinical-
decision agents." .

gov:CRR_CRD a gov:SectoralRegime ;
  rdfs:label "CRR/CRD"@en ;
  rdfs:comment "Capital Requirements Regulation 575/2013 plus Capital Requirements Directive
2013/36/EU." .

gov:IDD a gov:SectoralRegime ;
  rdfs:label "IDD"@en ;
  rdfs:comment "Insurance Distribution Directive 2016/97." .

gov:AWB a gov:SectoralRegime ;
  rdfs:label "Awb"@en ;
  rdfs:comment "Algemene wet bestuursrecht (Dutch General Administrative Law Act)." .

# -----
# Slot 17: gov:SupervisoryScope (ten Dutch supervisors plus EU bodies)
# Anchor: AI Act Art. 70; AP-RDI 7 November 2024
# -----

gov:SupervisoryScope a owl:Class ;
  rdfs:label "Supervisory Scope"@en ;
  rdfs:comment "Competent authority with defined supervisory mandate over an AI system or processing
activity." ;
  skos:note "Paper §4.3 slot 17; §5.2 ten-supervisor model." .

# Dutch supervisors (ten roles, nine entities)
gov:AFM a gov:SupervisoryScope ; rdfs:label "AFM"@en ; rdfs:comment "Autoriteit Financiële Markten." .
gov:DNB a gov:SupervisoryScope ; rdfs:label "DNB"@en ; rdfs:comment "De Nederlandsche Bank." .
gov:IGJ a gov:SupervisoryScope ; rdfs:label "IGJ"@en ; rdfs:comment "Inspectie Gezondheidszorg en
Jeugd." .
gov:AP a gov:SupervisoryScope ; rdfs:label "AP"@en ; rdfs:comment "Autoriteit Persoonsgegevens;
coordinating algorithm supervisor since February 2025." .
gov:RDI a gov:SupervisoryScope ; rdfs:label "RDI"@en ; rdfs:comment "Autoriteit Digitale
Infrastructuur." .
gov:NLA_SZW a gov:SupervisoryScope ; rdfs:label "Nederlandse Arbeidsinspectie"@en ; rdfs:comment
"Ministry of Social Affairs and Employment labour inspectorate." .
gov:NCSC a gov:SupervisoryScope ; rdfs:label "NCSC"@en ; rdfs:comment "Nationaal Cyber Security
Centrum." .
gov:ILT a gov:SupervisoryScope ; rdfs:label "ILT"@en ; rdfs:comment "Inspectie Leefomgeving en
Transport." .
gov:ACM_CM a gov:SupervisoryScope ; rdfs:label "ACM (consumer and markets)"@en ; rdfs:comment
"Autoriteit Consument en Markt, consumer-and-markets mandate." .
gov:ACM_Energy a gov:SupervisoryScope ; rdfs:label "ACM (energy)"@en ; rdfs:comment "ACM energy
division." .

# EU-level bodies
gov:EDPB a gov:SupervisoryScope ; rdfs:label "EDPB"@en ; rdfs:comment "European Data Protection
Board." .
gov:AIOffice a gov:SupervisoryScope ; rdfs:label "European AI Office"@en ; rdfs:comment "European
Commission AI Office." .

# -----
# Slot 18: gov:DelegationGrant + gov:RevocationEvent + gov:ValidityInterval + gov:AuthorityScope
# Anchor: GDPR Art. 28; AI Act Art. 25
# -----

gov:DelegationGrant a owl:Class ;
  rdfs:subClassOf prov:Delegation ;
  rdfs:label "Delegation Grant"@en ;
  rdfs:comment ""Temporal class with validity windows and revocation
semantics. Subclass of prov:Delegation extending it with explicit
validity intervals and authority scope. Required properties:
gov:hasValidityInterval, gov:hasAuthorityScope.""@en ;
  skos:note "Paper §4.3 slot 18; §6.1 revocation propagation; App F LTL formalisation." .

gov:RevocationEvent a owl:Class ;
  rdfs:subClassOf prov:Activity ;
  rdfs:label "Revocation Event"@en ;
  rdfs:comment "Event revoking a DelegationGrant. Triggers up to five effects per §6.1." .

gov:ValidityInterval a owl:Class ;
  rdfs:subClassOf time:ProperInterval ;
  rdfs:label "Validity Interval"@en ;
  rdfs:comment "Time-bounded validity for a delegation, role assignment, or policy version." .

gov:AuthorityScope a owl:Class ;
  rdfs:label "Authority Scope"@en ;
  rdfs:comment "The set of actions, resources, or data the delegated authority covers." .

```

```

gov:hasValidityInterval a owl:ObjectProperty ;
  rdfs:domain gov:DelegationGrant ;
  rdfs:range gov:ValidityInterval ;
  rdfs:label "has validity interval"@en .

gov:hasAuthorityScope a owl:ObjectProperty ;
  rdfs:domain gov:DelegationGrant ;
  rdfs:range gov:AuthorityScope ;
  rdfs:label "has authority scope"@en .

gov:isSubdelegationOf a owl:ObjectProperty ;
  rdfs:domain gov:DelegationGrant ;
  rdfs:range gov:DelegationGrant ;
  rdfs:label "is subdelegation of"@en ;
  rdfs:comment "Subdelegation chain. App A.9 sketch F6 specifies transitive revocation per GDPR Art.
28(2), (4)." .

# Five revocation effects (Section 6.1)
gov:BlocksNewWork a owl:Class ;
  rdfs:label "BlocksNewWork effect"@en ;
  rdfs:comment "Default effect of revocation: no new work accepted under revoked delegation." .

gov:CancelsQueuedWork a owl:Class ;
  rdfs:label "CancelsQueuedWork effect"@en ;
  rdfs:comment "Effect: queued (not yet executing) work under the delegation is cancelled." .

gov:StopsInflightWork a owl:Class ;
  rdfs:label "StopsInflightWork effect"@en ;
  rdfs:comment "Opt-in effect: in-progress tasks are stopped. Operationally risky default; specified
as opt-in per §6.1." .

gov:RequiresReview a owl:Class ;
  rdfs:label "RequiresReview effect"@en ;
  rdfs:comment "Effect: tasks completed recently under revoked authority must be reviewed before
acceptance downstream." .

gov:MarksArtifactsStale a owl:Class ;
  rdfs:label "MarksArtifactsStale effect"@en ;
  rdfs:comment "Default effect: artefacts produced under revoked authority marked stale; persistence
is a safety property (App F.3.5)." .

gov:executedUnderRevokedAuthority a owl:DatatypeProperty ;
  rdfs:domain gov:Decision ;
  rdfs:range xsd:boolean ;
  rdfs:label "executed under revoked authority"@en ;
  rdfs:comment "Default false. True requires explicit policy authorisation per §6.1." .

gov:cachedUnderRevokedAuthority a owl:DatatypeProperty ;
  rdfs:range xsd:boolean ;
  rdfs:label "cached under revoked authority"@en ;
  rdfs:comment "Artefact produced before revocation but accessed after; lawfully produced but may
need review." .

# -----
# Slot 19: gov:ProvenanceChain (integration point)
# Anchor: AI Act Art. 10; Art. 12
# -----
gov:ProvenanceChain a owl:Class ;
  rdfs:label "Provenance Chain"@en ;
  rdfs:comment ""Integration point reusing prov:wasDerivedFrom,
prov:wasInformedBy, prov:wasInfluencedBy. Counted as zero net new
classes under §2.5 counting convention (direct PROV-O reuse).""@en ;
  skos:note "Paper §4.3 slot 19." .

# -----
# Slot 20: gov:LogEvent (typed AI Act subclasses)
# Anchor: AI Act Art. 12(2)-(3); Art. 26(6); NIS2 Art. 23
# -----
gov:LogEvent a owl:Class ;
  rdfs:subClassOf prov:Activity ;
  rdfs:label "Log Event"@en ;
  rdfs:comment ""Typed log events at statutory-subclass granularity.
AI Act Art. 12(2) lists five categories implicitly; the MCS makes
them first-class subclasses. Generic logging is ubiquitous;
typed-to-statute subclasses are absent in surveyed ontologies.""@en ;
  skos:note "Paper §4.3 slot 20." .

gov:RiskIdentificationLogEvent a owl:Class ;
  rdfs:subClassOf gov:LogEvent ;
  rdfs:label "Risk Identification Log Event"@en ;
  rdfs:comment "AI Act Art. 12(2) (a): situations that may result in the AI system presenting a
risk." .

gov:PostMarketMonitoringLogEvent a owl:Class ;
  rdfs:subClassOf gov:LogEvent ;
  rdfs:label "Post-Market Monitoring Log Event"@en ;
  rdfs:comment "AI Act Art. 12(2) (b): post-market monitoring substantial modification events." .

gov:OperationMonitoringLogEvent a owl:Class ;
  rdfs:subClassOf gov:LogEvent ;
  rdfs:label "Operation Monitoring Log Event"@en ;
  rdfs:comment "AI Act Art. 12(2) (c): functional operation monitoring." .

gov:BiometricIdentificationLogEvent a owl:Class ;
  rdfs:subClassOf gov:LogEvent ;
  rdfs:label "Biometric Identification Log Event"@en ;
  rdfs:comment "AI Act Art. 12(3): mandatory fields for biometric identification systems." .

gov:IncidentLogEvent a owl:Class ;
  rdfs:subClassOf gov:LogEvent ;

```

```

rdfs:label "Incident Log Event"@en ;
rdfs:comment "AI Act Art. 73 plus NIS2 Art. 23 incident reporting events." .

# -----
# Slot 21: gov:MonitoringObservation
# Anchor: AI Act Art. 72; ECB Guide to Internal Models; SAFEST
# -----

gov:MonitoringObservation a owl:Class ;
rdfs:subClassOf prov:Entity ;
rdfs:label "Monitoring Observation"@en ;
rdfs:comment "Subclass of prov:Entity recording a metric observation in post-market monitoring
context." ;
skos:note "Paper §4.3 slot 21." .

gov:observedMetric a owl:DatatypeProperty ;
rdfs:domain gov:MonitoringObservation ;
rdfs:label "observed metric"@en .

gov:observedAtTime a owl:DatatypeProperty ;
rdfs:domain gov:MonitoringObservation ;
rdfs:range xsd:dateTime ;
rdfs:label "observed at time"@en .

gov:monitoringPurpose a owl:ObjectProperty ;
rdfs:domain gov:MonitoringObservation ;
rdfs:label "monitoring purpose"@en .

# -----
# Slot 22: gov:DriftIndicator (four subtypes)
# Anchor: AI Act Art. 72 continuous performance analysis; PCAOB AS 1105
# -----

gov:DriftIndicator a owl:Class ;
rdfs:label "Drift Indicator"@en ;
rdfs:comment "Indicator of drift in data, concept, model performance, or control." ;
skos:note "Paper §4.3 slot 22." .

gov:DataDrift a owl:Class ;
rdfs:subClassOf gov:DriftIndicator ;
rdfs:label "Data Drift"@en ;
rdfs:comment "Distributional shift in input data." .

gov:ConceptDrift a owl:Class ;
rdfs:subClassOf gov:DriftIndicator ;
rdfs:label "Concept Drift"@en ;
rdfs:comment "Shift in target-variable distribution or relationship between inputs and outputs." .

gov:ModelPerformanceDrift a owl:Class ;
rdfs:subClassOf gov:DriftIndicator ;
rdfs:label "Model Performance Drift"@en ;
rdfs:comment "Degradation in model performance metrics over time." .

gov:ControlDrift a owl:Class ;
rdfs:subClassOf gov:DriftIndicator ;
rdfs:label "Control Drift"@en ;
rdfs:comment "Erosion of operational controls (oversight cadence, intervention rates) over time."
.

gov:driftMagnitude a owl:DatatypeProperty ;
rdfs:domain gov:DriftIndicator ;
rdfs:range xsd:decimal ;
rdfs:label "drift magnitude"@en .

gov:driftDetectedAt a owl:DatatypeProperty ;
rdfs:domain gov:DriftIndicator ;
rdfs:range xsd:dateTime ;
rdfs:label "drift detected at"@en .

gov:triggersRefresh a owl:DatatypeProperty ;
rdfs:domain gov:DriftIndicator ;
rdfs:range xsd:boolean ;
rdfs:label "triggers refresh"@en ;
rdfs:comment "Whether the drift indicator triggers evidence-refresh per §6.2 Bayesian decay
framework." .

# -----
# Slot 23: gov:SupervisoryEvidenceArtifact + gov:Incident + gov:NonQualifyingADM (with five subtypes)
# Anchor: AI Act Art. 26(10); AI Act Art. 73; NIS2 Art. 23;
# GDPR Arts 5(1)(a), 6, 9, 25, 35; Art. 8 ECHR
# -----

gov:SupervisoryEvidenceArtifact a owl:Class ;
rdfs:subClassOf prov:Entity ;
rdfs:label "Supervisory Evidence Artifact"@en ;
rdfs:comment ""Evidence artefact with source hierarchy and probative weight.
Source hierarchy: SystemLogSource, HumanAttestationSource,
ExternalReportSource, ThirdPartyAuditSource (PCAOB AS 1105 illustrative
defaults, controller-configurable). Bayesian decay per §6.2.""@en ;
skos:note "Paper §4.3 slot 23 (one of three thematic classes); §6.2 evidence decay." .

gov:SystemLogSource a owl:Class ;
rdfs:label "System Log Source"@en ;
rdfs:comment "Evidence source: system-generated log. Default reliability weight 0.85 (App A.11)."
.

gov:HumanAttestationSource a owl:Class ;
rdfs:label "Human Attestation Source"@en ;
rdfs:comment "Evidence source: human attestation. Default reliability weight 0.60." .

gov:ExternalReportSource a owl:Class ;
rdfs:label "External Report Source"@en ;
rdfs:comment "Evidence source: external report. Default reliability weight 0.70." .

```

```

gov:ThirdPartyAuditSource a owl:Class ;
  rdfs:label "Third-Party Audit Source"@en ;
  rdfs:comment "Evidence source: third-party audit. Default reliability weight 0.90." .

gov:probativeWeight a owl:DatatypeProperty ;
  rdfs:domain gov:SupervisoryEvidenceArtifact ;
  rdfs:range xsd:decimal ;
  rdfs:label "probative weight"@en ;
  rdfs:comment ""Decimal in [0.0, 1.0]. Bayesian decay function with
  sectoral half-life priors per §6.2:
  ISAE 3402 12mo, ECB CRR/CRD 12mo, MDR 6mo, AI Act Art. 12 6mo (floor),
  AI Act Art. 72 12mo, NIS2 Art. 21 6mo, manual attestation 3mo.""@en .

gov:initialWeight a owl:DatatypeProperty ;
  rdfs:domain gov:SupervisoryEvidenceArtifact ;
  rdfs:range xsd:decimal ;
  rdfs:label "initial weight"@en .

gov:lastRefreshedAt a owl:DatatypeProperty ;
  rdfs:domain gov:SupervisoryEvidenceArtifact ;
  rdfs:range xsd:dateTime ;
  rdfs:label "last refreshed at"@en .

# Incident (AI Act Art. 73, NIS2 Art. 23 reporting)
gov:Incident a owl:Class ;
  rdfs:subClassOf prov:Activity ;
  rdfs:label "Serious Incident"@en ;
  rdfs:comment ""Serious incident under AI Act Art. 73 or NIS2 Art. 23.
  Significance assessed per F4 sketch (paper §5.3).""@en ;
  skos:note "Paper §4.3 slot 23 (second of three thematic classes); App A.7 sketch F4." .

gov:detectionTime a owl:DatatypeProperty ;
  rdfs:domain gov:Incident ;
  rdfs:range xsd:dateTime ;
  rdfs:label "detection time"@en .

gov:incidentWindow a owl:ObjectProperty ;
  rdfs:domain gov:Incident ;
  rdfs:range time:ProperInterval ;
  rdfs:label "incident window"@en ;
  rdfs:comment "Temporal window within which delegation events and memory writes must be preserved
  per NIS2 Art. 23(4); see App A.14, A.15." .

# NonQualifyingADM residual (paper §4.6 Finding 2)
gov:NonQualifyingADM a owl:Class ;
  rdfs:label "Non-Qualifying ADM"@en ;
  rdfs:comment ""Residual class for ADM enforcement decided on GDPR
  provisions other than Art. 22. Empirically anchored in App B.5
  (5 of 25 cases plus B22 SyRI on Art. 8 ECHR; six total). Sensitivity range
  [4%, 52%] under borderline-case flipping; Wilson 95% CI [11.5%, 43.4%].
  The category exists and is structurally significant; the proportion
  is empirically contingent.""@en ;
  skos:note "Paper §4.3 slot 23 (third of three thematic classes); §4.6 Finding 2." .

gov:Art5FairnessADM a owl:Class ;
  rdfs:subClassOf gov:NonQualifyingADM ;
  rdfs:label "Art. 5 Fairness ADM"@en ;
  rdfs:comment "ADM enforcement decided on GDPR Art. 5(1)(a) fairness ground. Empirical anchor: B20
  Belastingdienst Toeslagen." .

gov:Art6LawfulBasisADM a owl:Class ;
  rdfs:subClassOf gov:NonQualifyingADM ;
  rdfs:label "Art. 6 Lawful Basis ADM"@en ;
  rdfs:comment "ADM enforcement decided on GDPR Art. 6 lawful-basis ground. Empirical anchor: B18
  Clearview AI, B24 CNIL Clearview." .

gov:Art9SpecialCategoryADM a owl:Class ;
  rdfs:subClassOf gov:NonQualifyingADM ;
  rdfs:label "Art. 9 Special Category ADM"@en ;
  rdfs:comment "ADM enforcement decided on GDPR Art. 9 special-category-data ground. Empirical
  anchor: B19 live FR in schools." .

gov:Art25ByDesignADM a owl:Class ;
  rdfs:subClassOf gov:NonQualifyingADM ;
  rdfs:label "Art. 25 By-Design ADM"@en ;
  rdfs:comment "ADM enforcement decided on GDPR Art. 25 by-design ground. Empirical anchor: B17 IB
  grading." .

gov:Art35DPIAADM a owl:Class ;
  rdfs:subClassOf gov:NonQualifyingADM ;
  rdfs:label "Art. 35 DPIA ADM"@en ;
  rdfs:comment "ADM enforcement decided on GDPR Art. 35 DPIA-failure ground. Plus B22 SyRI decided
  on Art. 8 ECHR proportionality (cognate residual)." .

# =====
# Class-counting summary (Section 2.5 convention)
# =====
#
# Slots 1-23 declare 25 distinct row entries.
# Slot 3 (gov:Activity) is direct PROV-0 reuse with no new properties: 0 net new.
# Slot 19 (gov:ProvenanceChain) is direct PROV-0 reuse with no new properties: 0 net new.
# Slots 6-10 are DPV reuse: 0 net new each (5 rows total).
# Slot 23 counts three thematic classes (SupervisoryEvidenceArtifact, Incident,
# NonQualifyingADM) as one slot under the convention.
#
# 25 row entries minus 2 (slots 3 and 19) = 23 net new MCS classes.
#
# Under the alternative convention (each thematic class counted separately,
# subtype hierarchies flattened), the count ranges from approximately 30 to 43.
# Documented in paper §2.5; not adopted as the working figure.
#
# =====

```

```
# End of mcs_vocabulary.ttl
```

```
# =====
```

# Appendix H: SHACL profiles v1.2 (Turtle)

Source file: `mcs_profiles_v1_2.ttl`

```
# MCS Supervisory Profiles – v1.2 SHACL layer
#
# Companion artefact to:
#   "A Supervisory-Evidence Ontology for Agentic AI under EU Law:
#   Minimum Conceptual Set and Temporal Extension"
#   Jeroen Janssen, Apparens, April 2026
#
# This file declares the SHACL profile layer of the MCS:
# - Fourteen core temporal-extension shapes (paper §6.3): structural
#   validation of the temporal extension classes and properties, including
#   three SME-supporting shapes (HumanInterventionShape,
#   RevocationEventShape, LogEventSubclassGranularityShape) added in v0.5.1
#   to resolve previously dangling references in mcs_sme_profile_v0_2.ttl.
# - Profile A (AP-permissive, quantitative) and Profile B (CNIL-strict,
#   qualitative) implementations of the gov:Decision Article 22 scope test.
# - Profile C (UK ICO post-Brexit) and Profile D (Italian Garante post-Law
#   132/2025) ship as profile-class stubs only; no shapes attached pending
#   jurisdictional case-law review.
# - CrossProfileGovernanceShape and ProfileDeltaShape for surfacing
#   cross-profile divergences as governance artefacts.
#
# Total NodeShapes in this file: twenty-one (fourteen core temporal-extension
# shapes plus seven profile-layer shapes).
#
# Structural-completeness validation only. A sh:conforms=true result against
# any profile indicates structural conformance with that profile's shapes.
# It does not establish legal compliance with the AI Act, the GDPR, NIS2,
# DORA, the CRA, or any other instrument of Union or national law.
#
# © 2026 J. Janssen Apparens, CC BY 4.0

@prefix gov: <http://apparens.nl/mcs/gov#> .
@prefix sh: <http://www.w3.org/ns/shacl#> .
@prefix prov: <http://www.w3.org/ns/prov#> .
@prefix dpv: <https://w3id.org/dpv#> .
@prefix owl: <http://www.w3.org/2002/07/owl#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix skos: <http://www.w3.org/2004/02/skos/core#> .

# =====
# 1. gov:SupervisoryProfile abstract class
# =====

gov:SupervisoryProfile a owl:Class ;
  rdfs:label "Supervisory Profile"@en ;
  rdfs:comment """Abstract class parameter representing a jurisdiction-specific
  SHACL profile for Article 22 scope determination. A gov:Decision may be
  validated against one or more profiles; the delta between passing profiles
  is itself a governance artefact recorded as provenance.""@en ;
  skos:note "Paper reference: Section 5.5; Appendix A.3 dual-profile interaction." .

gov:hasAuthoritativeSource a owl:DatatypeProperty ;
  rdfs:domain gov:SupervisoryProfile ;
  rdfs:range xsd:anyURI ;
  rdfs:label "Authoritative source"@en ;
  rdfs:comment "URI of the DPA guidance, court decision, or academic source grounding the profile's
  interpretation." .

gov:hasDrawsStronglyOnMethod a owl:ObjectProperty ;
  rdfs:domain gov:SupervisoryProfile ;
  rdfs:range gov:DrawsStronglyOnMethod ;
  rdfs:label "drawsStronglyOn method"@en ;
  rdfs:comment "Which operationalisation of drawsStronglyOn the profile applies: quantitative (ε, M)
  or qualitative judicial indicia." .

gov:DrawsStronglyOnMethod a owl:Class ;
  rdfs:label "drawsStronglyOn Method"@en .

gov:QuantitativeFlipRate a gov:DrawsStronglyOnMethod ;
  rdfs:label "Quantitative flip rate"@en ;
  rdfs:comment "Counterfactual perturbation test: drawsStronglyOn holds iff flip rate at ε exceeds
  threshold M. Controller-side compliance instrument; may be rejected by supervisors." .

gov:QualitativeJudicialIndicia a gov:DrawsStronglyOnMethod ;
  rdfs:label "Qualitative judicial indicia"@en ;
  rdfs:comment "Aligned with Wiesbaden VG (January 2026) application: specialised scoring model,
  high refusal frequency given score (SCHUFA 'in almost all cases'), or absence of documented
  intervention authority. Judicially applied; controller-interpretable." .

# =====
# 2. Concrete profiles
# =====

gov:ProfileA a gov:SupervisoryProfile ;
  rdfs:label "Profile A (AP-permissive)"@en ;
  rdfs:comment "Autoriteit Persoonsgegevens reading: automated selection followed by meaningful
  human treatment falls outside Art. 22 engagement." ;
  gov:hasAuthoritativeSource
    "https://autoriteitpersoonsgegevens.nl/advies-art22-oct2024"^^xsd:anyURI ,
    "https://autoriteitpersoonsgegevens.nl/handvatten-juli2025"^^xsd:anyURI ;
  gov:hasDrawsStronglyOnMethod gov:QuantitativeFlipRate ;
  gov:quantitativeEpsilon "0.10"^^xsd:decimal ;
  gov:quantitativeM "0.50"^^xsd:decimal ;
  skos:note "Paper reference: Section 5.5 Profile A; Section 6.5 reference instantiation." .
```

```

gov:ProfileB a gov:SupervisoryProfile ;
  rdfs:label "Profile B (CNIL-strict)"@en ;
  rdfs:comment "CNIL, AEPD (March 2024), and several German state DPAs reading: score engages Art. 22(1) whenever downstream decision draws strongly on it, unless intervention meets a higher substantive bar. Wiesbaden VG (January 2026) applies qualitatively." ;
  gov:hasAuthoritativeSource
    "https://cnil.fr/ai-recommendations-2023-2025"^^xsd:anyURI ,
    "https://aepd.es/guidance-march-2024"^^xsd:anyURI ,
    "https://verwaltungsgericht-wiesbaden.de/january-2026"^^xsd:anyURI ;
  gov:hasDrawsStronglyOnMethod gov:QualitativeJudicialIndicia ;
  skos:note "Paper reference: Section 5.5 Profile B; Appendix B Finding 4 (drawsStronglyOn qualitative application)." .

gov:ProfileC a gov:SupervisoryProfile ;
  rdfs:label "Profile C (UK ICO post-Brexit, stub)"@en ;
  rdfs:comment "Stub for UK ICO alignment. 'Active, not a token gesture' standard for human review. Full population requires post-Brexit UK GDPR case law review. Not implemented at v1.2 release." ;
  gov:hasAuthoritativeSource
    "https://ico.org.uk/ai-guidance"^^xsd:anyURI ;
  gov:hasDrawsStronglyOnMethod gov:QualitativeJudicialIndicia .

gov:ProfileD a gov:SupervisoryProfile ;
  rdfs:label "Profile D (Italian Garante post-Law 132/2025, stub)"@en ;
  rdfs:comment "Stub for Italian Garante alignment under Law 132/2025 sector-specific obligations (healthcare, labour, justice, public administration). Full population requires sectoral variant mapping in mcs_framework_bridges.ttl. Not implemented at v1.2 release." ;
  gov:hasAuthoritativeSource
    "https://gdpd.it/ai-guidance"^^xsd:anyURI ,
    "https://normattiva.it/uri-res/N2Ls?urn:nir:stato:legge:2025-09-25;132"^^xsd:anyURI ;
  gov:hasDrawsStronglyOnMethod gov:QualitativeJudicialIndicia .

# =====
# 3. Profile A shapes: AP-permissive + quantitative drawsStronglyOn
# =====

gov:ProfileA_Art22ScopeShape a sh:NodeShape ;
  rdfs:label "Profile A Art. 22 scope shape"@en ;
  sh:targetClass gov:Score ;
  sh:sparql [
    sh:prefixes [
      sh:declare [ sh:prefix "gov" ; sh:namespace "http://apparens.nl/mcs/gov#"^^xsd:anyURI ]
    ] ;
    sh:select """
      SELECT $this WHERE {
        $this a gov:Score .
        ?d gov:drawsStronglyOn $this ;
          gov:hasLegalEffect true .
        FILTER NOT EXISTS {
          ?d a gov:AssistedDecisionWithMeaningfulHumanIntervention .
        }
      }
    """ ;
    sh:message "Profile A (AP-permissive): Score engages Art. 22(1) via downstream legal-effect decision drawing strongly on it, unless the decision is classified as AssistedDecisionWithMeaningfulHumanIntervention." ;
    sh:severity sh:Info
  ] .

gov:ProfileA_DrawsStronglyOnShape a sh:NodeShape ;
  rdfs:label "Profile A drawsStronglyOn quantitative shape"@en ;
  sh:targetSubjectsOf gov:drawsStronglyOn ;
  sh:property [
    sh:path gov:flipRateAtEpsilon10 ;
    sh:datatype xsd:decimal ;
    sh:minCount 1 ;
    sh:message "Profile A requires gov:flipRateAtEpsilon10 populated on any gov:Decision invoking drawsStronglyOn. Flip rate ≥ 0.50 at ε = 10% of population range satisfies the quantitative test."
  ] ;
  sh:property [
    sh:path gov:flipRateAtEpsilon10 ;
    sh:minInclusive "0.50"^^xsd:decimal ;
    sh:severity sh:Info ;
    sh:message "Profile A quantitative test satisfied: flip rate at ε=10% ≥ 50%, drawsStronglyOn holds."
  ] .

# =====
# 4. Profile B shapes: CNIL-strict + qualitative judicial indicia
# =====

gov:ProfileB_Art22ScopeShape a sh:NodeShape ;
  rdfs:label "Profile B Art. 22 scope shape"@en ;
  sh:targetClass gov:Score ;
  sh:sparql [
    sh:prefixes [
      sh:declare [ sh:prefix "gov" ; sh:namespace "http://apparens.nl/mcs/gov#"^^xsd:anyURI ]
    ] ;
    sh:select """
      SELECT $this WHERE {
        $this a gov:Score .
        ?d gov:drawsStronglyOn $this ;
          gov:hasLegalEffect true .
        FILTER NOT EXISTS {
          ?d gov:hasHumanIntervention ?hi .
          ?hi gov:hasAuthorityToDeviate true ;
            gov:hasMarginOfDiscretion true ;
            gov:hasUnderstandingOfLogic true ;
            gov:consideredAllRelevantData true ;
            gov:authorityActuallyExercised true .
        }
      }
    """ ;
  ] .

```

```

    sh:message "Profile B (CNIL-strict): Score engages Art. 22(1) via downstream legal-effect
decision drawing strongly on it, unless intervention meets all five conjunctive properties including
gov:authorityActuallyExercised (not merely gov:hasAuthorityToDeviate). The fifth property responds to
Uber/Ola Amsterdam Court of Appeal finding that Krakow review was 'not much more than a purely
symbolic act'." ;
    sh:severity sh:Info
] .

gov:ProfileB_DrawsStronglyOnShape a sh:NodeShape ;
rdfs:label "Profile B drawsStronglyOn qualitative shape"@en ;
sh:targetSubjectsOf gov:drawsStronglyOn ;
sh:or (
[
    sh:property [
        sh:path gov:usesSpecialisedScoringModel ;
        sh:hasValue true ;
        sh:message "Profile B qualitative indicator 1 satisfied: use of specialised scoring
model (Wiesbaden VG January 2026)."
```

```

rdfs:range gov:SupervisoryProfile ;
rdfs:label "validated against profile"@en .

gov:passesProfile a owl:ObjectProperty ;
rdfs:domain gov:Decision ;
rdfs:range gov:SupervisoryProfile ;
rdfs:label "passes profile"@en ;
rdfs:comment "Decision passes the SHACL validation shapes associated with the named profile." .

gov:hasProfileDelta a owl:ObjectProperty ;
rdfs:domain gov:Decision ;
rdfs:range gov:ProfileDelta ;
rdfs:label "has profile delta"@en .

gov:ProfileDelta a owl:Class ;
rdfs:label "Profile Delta"@en ;
rdfs:comment "First-class governance artefact recording the divergence between supervisory
profiles for a given decision." .

gov:authorityActuallyExercised a owl:DatatypeProperty ;
rdfs:domain gov:HumanIntervention ;
rdfs:range xsd:boolean ;
rdfs:label "authority actually exercised"@en ;
rdfs:comment "True iff the human intervener actually exercised authority to deviate (demonstrated
by deviation rate > declared threshold over the intervention population), as distinct from merely
possessing the authority. Responds to Uber/Ola Hof Amsterdam 'symbolic act' finding
(ECLI:NL:GHAMS:2023:793)." .

gov:flipRateAtEpsilon10 a owl:DatatypeProperty ;
rdfs:domain gov:Decision ;
rdfs:range xsd:decimal ;
rdfs:label "flip rate at epsilon 10%"@en ;
rdfs:comment "Profile A quantitative measure: proportion of perturbation replicas in which the
decision outcome flips when the score is perturbed by  $\epsilon = 10\%$  of population range." .

gov:flipRateAtEpsilon20 a owl:DatatypeProperty ;
rdfs:domain gov:Decision ;
rdfs:range xsd:decimal ;
rdfs:label "flip rate at epsilon 20%"@en ;
rdfs:comment "Alternative quantitative measure at  $\epsilon = 20\%$  of population range. Retained for
controllers who wish to populate both profiles quantitatively; Profile B normatively uses qualitative
indicia." .

gov:usesSpecialisedScoringModel a owl:DatatypeProperty ;
rdfs:domain gov:Decision ;
rdfs:range xsd:boolean ;
rdfs:label "uses specialised scoring model"@en ;
rdfs:comment "Profile B qualitative indicator 1. Aligned with Wiesbaden VG (January 2026): the
existence of a specialised scoring model is treated as evidence of strong reliance without requiring
quantification." .

gov:refusalFrequencyGivenScore a owl:DatatypeProperty ;
rdfs:domain gov:Decision ;
rdfs:range xsd:decimal ;
rdfs:label "refusal frequency given score"@en ;
rdfs:comment "Profile B qualitative indicator 2. Aligned with SCHUFA paragraph 48 'in almost all
cases' factual threshold. Conditional probability of refusal given insufficient score in the
controller's historical decision population. Threshold 0.90 operationalises 'in almost all cases'
pending jurisprudential refinement." .

gov:hasDocumentedIntervention a owl:DatatypeProperty ;
rdfs:domain gov:Decision ;
rdfs:range xsd:boolean ;
rdfs:label "has documented intervention"@en ;
rdfs:comment "Profile B qualitative indicator 3. False iff the decision lacks any documented human
intervention with the five Profile B conjunctive properties. Aligned with Dun and Bradstreet paragraph
40 'lacked any manual oversight' factual finding." .

gov:quantitativeEpsilon a owl:DatatypeProperty ;
rdfs:domain gov:SupervisoryProfile ;
rdfs:range xsd:decimal ;
rdfs:label "quantitative epsilon"@en ;
rdfs:comment "Perturbation magnitude  $\epsilon$ , expressed as proportion of population range, used in
quantitative drawsStronglyOn operationalisation." .

gov:quantitativeM a owl:DatatypeProperty ;
rdfs:domain gov:SupervisoryProfile ;
rdfs:range xsd:decimal ;
rdfs:label "quantitative M"@en ;
rdfs:comment "Flip rate threshold M used in quantitative drawsStronglyOn operationalisation." .

# =====
# Core temporal-extension shapes (eleven shapes; paper §6.3)
# =====
# These shapes operationalise the structural-completeness checks
# for the temporal extension classes. They fire as structural
# completeness validation, not legal-compliance validation.

gov:DelegationGrantShape a sh:NodeShape ;
sh:targetClass gov:DelegationGrant ;
rdfs:label "Delegation Grant Shape"@en ;
rdfs:comment "Validates that a DelegationGrant declares both a validity interval and an authority
scope (paper §6.3, slot 18)." ;
sh:property [
  sh:path gov:hasValidityInterval ;
  sh:minCount 1 ;
  sh:class gov:ValidityInterval ;
  sh:message "Every DelegationGrant must declare a validity interval (gov:hasValidityInterval)."
] ;
sh:property [
  sh:path gov:hasAuthorityScope ;
  sh:minCount 1 ;
  sh:class gov:AuthorityScope ;

```

```

    sh:message "Every DelegationGrant must declare an authority scope (gov:hasAuthorityScope)."
  ] .

gov:PolicyVersionShape a sh:NodeShape ;
  sh:targetClass gov:PolicyVersion ;
  rdfs:label "Policy Version Shape"@en ;
  rdfs:comment "Validates that a PolicyVersion has a validity interval and that supersession does
not produce overlapping intervals with its predecessor (paper §6.3, slot 13).";
  sh:property [
    sh:path gov:policyValidityInterval ;
    sh:minCount 1 ;
    sh:class prov:Entity ;
    sh:message "Every PolicyVersion must declare a validity interval
(gov:policyValidityInterval).";
  ] ;
  sh:property [
    sh:path gov:supersedesPolicy ;
    sh:maxCount 1 ;
    sh:class gov:PolicyVersion ;
    sh:message "A PolicyVersion supersedes at most one predecessor; chained supersession is
recorded transitively, not via multiple direct edges."
  ] .

gov:EvidenceDecayShape a sh:NodeShape ;
  sh:targetClass gov:SupervisoryEvidenceArtifact ;
  rdfs:label "Evidence Decay Shape"@en ;
  rdfs:comment "Validates that a SupervisoryEvidenceArtifact carries the timestamp required to
compute Bayesian decay per declared sectoral priors (paper §6.2, §6.3, slot 23).";
  sh:property [
    sh:path gov:lastRefreshedAt ;
    sh:minCount 1 ;
    sh:datatype xsd:dateTime ;
    sh:message "Every SupervisoryEvidenceArtifact must declare gov:lastRefreshedAt; without it,
decay function cannot be evaluated."
  ] ;
  sh:property [
    sh:path gov:probativeWeight ;
    sh:datatype xsd:decimal ;
    sh:minInclusive "0.0"^^xsd:decimal ;
    sh:maxInclusive "1.0"^^xsd:decimal ;
    sh:message "gov:probativeWeight, where present, must be in [0.0, 1.0]."
  ] .

gov:EvidenceQualityShape a sh:NodeShape ;
  sh:targetClass gov:SupervisoryEvidenceArtifact ;
  rdfs:label "Evidence Quality Shape"@en ;
  rdfs:comment "Validates that a SupervisoryEvidenceArtifact declares its source type, enabling
PCAOB AS 1105 reliability weighting (paper §6.3, App A.11 sketch F8).";
  sh:property [
    sh:path prov:wasGeneratedBy ;
    sh:minCount 1 ;
    sh:message "Every SupervisoryEvidenceArtifact must be linked to its generating activity via
prov:wasGeneratedBy."
  ] .

gov:MonitoringObservationShape a sh:NodeShape ;
  sh:targetClass gov:MonitoringObservation ;
  rdfs:label "Monitoring Observation Shape"@en ;
  rdfs:comment "Validates that a MonitoringObservation carries the metric, timestamp, and purpose
required for AI Act Art. 72 post-market monitoring (paper §6.3, slot 21).";
  sh:property [
    sh:path gov:observedMetric ;
    sh:minCount 1 ;
    sh:message "Every MonitoringObservation must declare a metric (gov:observedMetric).";
  ] ;
  sh:property [
    sh:path gov:observedAtTime ;
    sh:minCount 1 ;
    sh:datatype xsd:dateTime ;
    sh:message "Every MonitoringObservation must declare gov:observedAtTime."
  ] .

gov:DriftIndicatorShape a sh:NodeShape ;
  sh:targetClass gov:DriftIndicator ;
  rdfs:label "Drift Indicator Shape"@en ;
  rdfs:comment "Validates that a DriftIndicator carries magnitude, detection time, and refresh-
trigger signal (paper §6.3, slot 22).";
  sh:property [
    sh:path gov:driftMagnitude ;
    sh:minCount 1 ;
    sh:datatype xsd:decimal ;
    sh:message "Every DriftIndicator must declare gov:driftMagnitude."
  ] ;
  sh:property [
    sh:path gov:driftDetectedAt ;
    sh:minCount 1 ;
    sh:datatype xsd:dateTime ;
    sh:message "Every DriftIndicator must declare gov:driftDetectedAt."
  ] .

gov:AgentActionShape a sh:NodeShape ;
  sh:targetClass gov:Activity ;
  rdfs:label "Agent Action Shape"@en ;
  rdfs:comment "Validates that an Activity tied to a Decision pipeline declares lawful basis,
purpose, and (where applicable) data subject reference, per the integration pattern in paper §4.2." ;
  sh:property [
    sh:path dpv:hasLegalBasis ;
    sh:minCount 1 ;
    sh:class dpv:LegalBasis ;
    sh:message "Activity contributing to a Decision must declare a dpv:hasLegalBasis under GDPR
Art. 6(1).";
  ] ;
  sh:property [

```

```

    sh:path dpv:hasPurpose ;
    sh:minCount 1 ;
    sh:class dpv:Purpose ;
    sh:message "Activity contributing to a Decision must declare a dpv:hasPurpose under GDPR Art.
5(1)(b)."
] .

gov:RoleAssignmentShape a sh:NodeShape ;
sh:targetClass gov:RoleAssignment ;
rdfs:label "Role Assignment Shape"@en ;
rdfs:comment "Validates that a RoleAssignment is time-bounded so that role ownership can be
reconstructed at any past moment (paper §4.2, slot 2).";
sh:property [
  sh:path gov:hasAgent ;
  sh:minCount 1 ;
  sh:class gov:Agent ;
  sh:message "Every RoleAssignment must point to a gov:Agent (gov:hasAgent).";
];
sh:property [
  sh:path gov:hasRole ;
  sh:minCount 1 ;
  sh:class gov:LegalRole ;
  sh:message "Every RoleAssignment must point to a gov:LegalRole (gov:hasRole).";
];
sh:property [
  sh:path gov:assignmentInterval ;
  sh:minCount 1 ;
  sh:message "Every RoleAssignment must declare an assignmentInterval; untimed role attribution
defeats reconstruction."
] .

gov:BiometricLogShape a sh:NodeShape ;
sh:targetClass gov:BiometricIdentificationLogEvent ;
rdfs:label "Biometric Log Shape"@en ;
rdfs:comment "Validates the AI Act Art. 12(3) mandatory fields for biometric identification
systems: start-and-end timestamps, reference database identification, search-result identification,
and oversight-natural-person identification (paper §6.3, slot 20).";
sh:property [
  sh:path prov:startedAtTime ;
  sh:minCount 1 ;
  sh:datatype xsd:dateTime ;
  sh:message "AI Act Art. 12(3): biometric log must record start time."
];
sh:property [
  sh:path prov:endedAtTime ;
  sh:minCount 1 ;
  sh:datatype xsd:dateTime ;
  sh:message "AI Act Art. 12(3): biometric log must record end time."
];
sh:property [
  sh:path prov:wasAssociatedWith ;
  sh:minCount 1 ;
  sh:class gov:Agent ;
  sh:message "AI Act Art. 12(3): biometric log must identify the natural person carrying out the
verification."
] .

gov:DecisionShape a sh:NodeShape ;
sh:targetClass gov:Decision ;
rdfs:label "Decision Shape"@en ;
rdfs:comment "Validates that a Decision is well-formed under the three-way taxonomy: exactly one
classification, decision time recorded, and (where the Decision is potentially Art. 22-engaging)
hasLegalEffect declared (paper §4.3, slot 11).";
sh:property [
  sh:path rdf:type ;
  sh:hasValue gov:Decision ;
  sh:message "Decision shape applies to instances explicitly typed as gov:Decision."
];
sh:or (
  [ sh:property [ sh:path [ sh:inversePath rdf:type ] ; sh:hasValue gov:SolelyAutomatedDecision
] ]
  [ sh:property [ sh:path [ sh:inversePath rdf:type ] ; sh:hasValue
gov:AssistedDecisionWithMeaningfulHumanIntervention ] ]
  [ sh:property [ sh:path [ sh:inversePath rdf:type ] ; sh:hasValue gov:NonAutomatedDecision ] ]
);
sh:property [
  sh:path gov:hasLegalEffect ;
  sh:datatype xsd:boolean ;
  sh:message "Decision must declare gov:hasLegalEffect (boolean) so Art. 22(1) qualifier is
structurally evaluable."
] .

gov:SchufaUpstreamShape a sh:NodeShape ;
sh:targetClass gov:Score ;
rdfs:label "SCHUFA Upstream Shape"@en ;
rdfs:comment "Validates that gov:Score carries the predicates required to evaluate drawsStronglyOn
under both Profile A (quantitative) and Profile B (qualitative) operationalisations (paper §6.4).";
sh:or (
  # Profile A path: quantitative flip rate present
  [ sh:property [
    sh:path gov:flipRateAtEpsilon10 ;
    sh:minCount 1 ;
    sh:datatype xsd:decimal
  ] ]
  # Profile B path: at least one qualitative indicium evaluable
  [ sh:property [
    sh:path gov:usesSpecialisedScoringModel ;
    sh:minCount 1 ;
    sh:datatype xsd:boolean
  ] ]
);
rdfs:comment "A Score consumed by an Art. 22-engaging Decision must support evaluation under at
least one drawsStronglyOn operationalisation. Failure indicates the controller has not populated

```

```

either Profile A's quantitative property or any of Profile B's three qualitative indicia." .

# =====
# 5. SME-supporting core shapes added in v0.5.1
# =====
# Three shapes added to resolve dangling references previously present in
# mcs_sme_profile_v0_2.ttl (gov:HumanInterventionShape,
# gov:RevocationEventShape, gov:LogEventSubclassGranularityShape).
# These shapes operationalise structural-completeness checks for the
# slot 12 (HumanIntervention), §6.1 (RevocationEvent), and slot 20
# (typed LogEvent) constructs respectively.

gov:HumanInterventionShape a sh:NodeShape ;
  sh:targetClass gov:HumanIntervention ;
  rdfs:label "Human Intervention Shape"@en ;
  rdfs:comment "Validates that a HumanIntervention carries the six conjunctive properties of the
slot 12 meaningful-intervention schema (paper §4.3 slot 12; §4.6 Finding 3). Profile B in particular
requires gov:authorityActuallyExercised=true; this shape only checks structural presence of the six
properties, not their truth values." ;
  sh:property [
    sh:path gov:hasAuthorityToDeviate ;
    sh:minCount 1 ;
    sh:datatype xsd:boolean ;
    sh:message "Property 1 of slot 12: gov:hasAuthorityToDeviate must be declared (Boolean)."
```

# Appendix I: SME proportionality profile v0.2 (Turtle)

Source file: `mcs_sme_profile_v0_2.ttl`

```
# MCS SME Proportionality Profile - v0.1
# Size-based proportionality profile for AI Act high-risk controllers
# meeting the EU SME definition (Recommendation 2003/361/EC).
#
# Complementary to mcs_profiles_v1_2.ttl dual-jurisdictional profiles.
# SME proportionality operates on a separate axis (size, not jurisdiction).
#
# Structural-completeness validation only. An SME running under this
# profile remains subject to the GDPR, the AI Act, and NIS2 at the
# same substantive standard as a larger controller. This file reduces
# the evidentiary machinery that the MCS specifies as the default,
# not the underlying legal duty.
#
# © 2026 J. Janssen Apparens, CC BY 4.0 (more permissive than paper
# licence to allow downstream adaptation by SME-sector bodies and
# data-subject representation organisations).

@prefix gov: <http://apparens.nl/mcs/gov#> .
@prefix sh: <http://www.w3.org/ns/shacl#> .
@prefix prov: <http://www.w3.org/ns/prov#> .
@prefix owl: <http://www.w3.org/2002/07/owl#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix skos: <http://www.w3.org/2004/02/skos/core#> .

# =====
# 1. Eligibility criterion class
# =====

gov:EligibilityCriterion a owl:Class ;
  rdfs:label "Eligibility Criterion"@en ;
  rdfs:comment "Abstract class for size-based or sector-based eligibility criteria that parameterise
proportionality profiles." .

gov:EC_SME_Definition_2003_361 a gov:EligibilityCriterion ;
  rdfs:label "EU SME Definition (Recommendation 2003/361/EC)"@en ;
  gov:maxEmployeeCount "250"^^xsd:integer ;
  gov:maxAnnualTurnover "50000000"^^xsd:decimal ;
  gov:maxBalanceSheetTotal "43000000"^^xsd:decimal ;
  gov:eligibilityConjunction gov:EmployeeCountAndEitherFinancialTest ;
  gov:eligibilityReassessmentCadence "PIY"^^xsd:duration ;
  gov:authoritativeSource "https://eur-lex.europa.eu/legal-content/EN/TXT/?
uri=celex:32003H0361"^^xsd:anyURI .

gov:EmployeeCountAndEitherFinancialTest a owl:Class ;
  rdfs:comment "EU SME definition eligibility: employee count below threshold AND (turnover below
threshold OR balance sheet below threshold)." .

# Member State stricter threshold example (for reference; not activated by default)
gov:EC_MicroSME_Definition a gov:EligibilityCriterion ;
  rdfs:label "Micro SME Definition (under EU Rec 2003/361/EC)"@en ;
  rdfs:comment "Stricter sub-category. Member States applying micro-SME-specific measures reference
this instead." ;
  gov:maxEmployeeCount "10"^^xsd:integer ;
  gov:maxAnnualTurnover "2000000"^^xsd:decimal ;
  gov:maxBalanceSheetTotal "2000000"^^xsd:decimal .

# =====
# 2. SME proportionality profile as gov:SupervisoryProfile instance
# =====

gov:SMEEExemptionProfile a gov:SupervisoryProfile ;
  rdfs:label "SME Proportionality Profile"@en ;
  rdfs:comment "Size-based proportionality profile reducing evidentiary machinery for SMEs meeting
EU Recommendation 2003/361/EC eligibility. Does not reduce substantive legal obligation." ;
  gov:hasEligibilityCriterion gov:EC_SME_Definition_2003_361 ;
  gov:hasSubstitutionMechanism gov:StructuredProseNarrative ;
  gov:hasExemptedShapeSet gov:SMEEExemptedShapes_v0_1 ;
  gov:hasRetainedShapeSet gov:SMERetainedShapes_v0_1 ;
  gov:compatibleWith gov:ProfileA, gov:ProfileB ;
  skos:note "Paper reference: §5.6." .

gov:StructuredProseNarrative a owl:Class ;
  rdfs:label "Structured Prose Narrative"@en ;
  rdfs:comment "Substitution mechanism: prose narrative with declared structure (what evidence would
have populated the shape, why disproportionate, what alternative mechanism applies)." .

# =====
# 3. Retained shape set (in-scope for SMEs)
# =====

gov:SMERetainedShapes_v0_1 a gov:ShapeSet ;
  rdfs:label "SME Retained Shapes v0.1"@en ;
  rdfs:comment "SHACL shapes that remain in full validation scope for SMEs running under the SME
proportionality profile. Legacy URI retained: gov:SMEEExemptionProfile." ;
  gov:containsShape
    gov:DecisionShape ,
    gov:HumanInterventionShape ,
    gov:PolicyVersionShape ,
    gov:DelegationGrantShape ,
    gov:RevocationEventShape ,
    gov:ProfileA_Art22ScopeShape ,
    gov:ProfileB_Art22ScopeShape .

# =====
# 4. Exempted shape set (prose-substitutable for SMEs)
```

```

# =====
gov:SMEExemptedShapes_v0_1 a gov:ShapeSet ;
rdfs:label "SME Exempted Shapes v0.1"@en ;
rdfs:comment "SHACL shapes that an SME may substitute with structured prose narrative." ;
gov:containsShape
  gov:LogEventSubclassGranularityShape ,
  gov:MonitoringObservationShape ,
  gov:DriftIndicatorShape ,
  gov:EvidenceDecayShape ,
  gov:CrossProfileGovernanceShape .

# =====
# 5. Proportionality marker shape
# =====

gov:SMEExemptionMarker a sh:NodeShape ;
rdfs:label "SME Proportionality Marker Shape"@en ;
sh:targetSubjectsOf gov:validatedAgainstProfile ;
sh:sparql [
  sh:prefixes [
    sh:declare [ sh:prefix "gov" ; sh:namespace "http://apparens.nl/mcs/gov#"^^xsd:anyURI ]
  ] ;
  sh:select """
    SELECT $this ?exemptedShape WHERE {
      $this gov:validatedAgainstProfile gov:SMEExemptionProfile .
      gov:SMEExemptionProfile gov:hasExemptedShapeSet ?shapeSet .
      ?shapeSet gov:containsShape ?exemptedShape .
      FILTER NOT EXISTS {
        $this gov:hasProseSubstitution ?ps .
        ?ps gov:substitutesFor ?exemptedShape .
      }
    }
    """ ;
  sh:message "Subject validated against SME proportionality profile must supply
gov:hasProseSubstitution for every exempted shape. Missing substitution narrative makes the
proportionality profile structurally invalid." ;
  sh:severity sh:Violation
] .

# =====
# 6. Prose substitution structure
# =====

gov:ProseSubstitution a owl:Class ;
rdfs:subClassOf prov:Entity ;
rdfs:label "Prose Substitution"@en ;
rdfs:comment "Structured prose narrative substituting for a specific exempted SHACL shape." .

gov:ProseSubstitutionShape a sh:NodeShape ;
sh:targetClass gov:ProseSubstitution ;
sh:property [
  sh:path gov:substitutesFor ;
  sh:minCount 1 ;
  sh:message "Every ProseSubstitution must identify which shape it substitutes for."
] ;
sh:property [
  sh:path gov:wouldHavePopulatedWith ;
  sh:minCount 1 ;
  sh:datatype xsd:string ;
  sh:message "Narrative must declare what evidence would have populated the shape if non-
exempt."
] ;
sh:property [
  sh:path gov:disproportionalityJustification ;
  sh:minCount 1 ;
  sh:datatype xsd:string ;
  sh:message "Narrative must declare why SHACL-validated evidence production is disproportionate
at the controller's operational scale."
] ;
sh:property [
  sh:path gov:alternativeAccountabilityMechanism ;
  sh:minCount 1 ;
  sh:datatype xsd:string ;
  sh:message "Narrative must declare what alternative accountability mechanism the controller
provides in place of SHACL-validated evidence (for example, a documented manual review protocol)."
] ;
sh:property [
  sh:path prov:generatedAtTime ;
  sh:minCount 1 ;
  sh:datatype xsd:dateTime ;
  sh:message "Narrative must be time-stamped."
] .

# =====
# 7. Supervisor-initiated proportionality-treatment suspension
# =====

gov:ExemptionSuspension a owl:Class ;
rdfs:subClassOf prov:Activity ;
rdfs:label "Proportionality Treatment Suspension"@en ;
rdfs:comment "Supervisor-initiated suspension of SME proportionality treatment for a specified
dimension, scope, and window. During the suspension, the suspended shape(s) must be populated in full
SHACL-validated form. Outside the window, the proportionality arrangement resumes automatically." .

gov:ExemptionSuspensionShape a sh:NodeShape ;
sh:targetClass gov:ExemptionSuspension ;
sh:property [
  sh:path gov:suspendsShape ;
  sh:minCount 1 ;
  sh:message "Suspension must identify which exempted shape(s) the suspension applies to."
] ;
sh:property [

```

```

        sh:path prov:startedAtTime ;
        sh:minCount 1 ;
        sh:datatype xsd:dateTime ;
        sh:message "Suspension must have a start time."
    ] ;
    sh:property [
        sh:path prov:endedAtTime ;
        sh:minCount 1 ;
        sh:datatype xsd:dateTime ;
        sh:message "Suspension must have an end time. Permanent revocation of SME eligibility is a
different mechanism; this class is for time-bounded investigation suspensions."
    ] ;
    sh:property [
        sh:path gov:suspendingAuthority ;
        sh:minCount 1 ;
        sh:class gov:SupervisoryScope ;
        sh:message "Suspension must identify the supervisory authority issuing it."
    ] ;
    sh:property [
        sh:path gov:suspensionScope ;
        sh:minCount 1 ;
        sh:datatype xsd:string ;
        sh:message "Suspension must declare its scope (for example, investigation file reference,
decision subset)."
```

].

```

# =====
# 8. Properties
# =====

gov:hasEligibilityCriterion a owl:ObjectProperty ;
    rdfs:domain gov:SupervisoryProfile ;
    rdfs:range gov:EligibilityCriterion .

gov:hasSubstitutionMechanism a owl:ObjectProperty ;
    rdfs:domain gov:SupervisoryProfile ;
    rdfs:range owl:Class .

gov:hasExemptedShapeSet a owl:ObjectProperty ;
    rdfs:domain gov:SupervisoryProfile ;
    rdfs:range gov:ShapeSet .

gov:hasRetainedShapeSet a owl:ObjectProperty ;
    rdfs:domain gov:SupervisoryProfile ;
    rdfs:range gov:ShapeSet .

gov:containsShape a owl:ObjectProperty ;
    rdfs:domain gov:ShapeSet ;
    rdfs:range sh:NodeShape .

gov:hasProseSubstitution a owl:ObjectProperty ;
    rdfs:range gov:ProseSubstitution .

gov:substitutesFor a owl:ObjectProperty ;
    rdfs:domain gov:ProseSubstitution ;
    rdfs:range sh:NodeShape .

gov:wouldHavePopulatedWith a owl:DatatypeProperty ;
    rdfs:domain gov:ProseSubstitution ;
    rdfs:range xsd:string .

gov:disproportionalityJustification a owl:DatatypeProperty ;
    rdfs:domain gov:ProseSubstitution ;
    rdfs:range xsd:string .

gov:alternativeAccountabilityMechanism a owl:DatatypeProperty ;
    rdfs:domain gov:ProseSubstitution ;
    rdfs:range xsd:string .

gov:maxEmployeeCount a owl:DatatypeProperty ;
    rdfs:domain gov:EligibilityCriterion ;
    rdfs:range xsd:integer .

gov:maxAnnualTurnover a owl:DatatypeProperty ;
    rdfs:domain gov:EligibilityCriterion ;
    rdfs:range xsd:decimal .

gov:maxBalanceSheetTotal a owl:DatatypeProperty ;
    rdfs:domain gov:EligibilityCriterion ;
    rdfs:range xsd:decimal .

gov:eligibilityConjunction a owl:ObjectProperty ;
    rdfs:domain gov:EligibilityCriterion ;
    rdfs:range owl:Class .

gov:eligibilityReassessmentCadence a owl:DatatypeProperty ;
    rdfs:domain gov:EligibilityCriterion ;
    rdfs:range xsd:duration .

gov:authoritativeSource a owl:DatatypeProperty ;
    rdfs:range xsd:anyURI .

gov:suspendsShape a owl:ObjectProperty ;
    rdfs:domain gov:ExemptionSuspension ;
    rdfs:range sh:NodeShape .

gov:suspendingAuthority a owl:ObjectProperty ;
    rdfs:domain gov:ExemptionSuspension ;
    rdfs:range gov:SupervisoryScope .

gov:suspensionScope a owl:DatatypeProperty ;
    rdfs:domain gov:ExemptionSuspension ;
    rdfs:range xsd:string .

```

```
gov:compatibleWith a owl:ObjectProperty ;
  rdfs:domain gov:SupervisoryProfile ;
  rdfs:range gov:SupervisoryProfile ;
  rdfs:comment "The SME proportionality profile is compatible with both Profile A and Profile B
jurisdictional choices. A controller runs both axes simultaneously: one jurisdictional profile and
optionally the SME axis." .

gov:ShapeSet a owl:Class ;
  rdfs:label "Shape Set"@en .
```

## Appendix J: Revocation model v0.2 (TLA+)

Source file: `tla_model_checking/mcs_revocation_model_v0.2.tla`

```
---- MODULE McsRevocationModel ----
\*
\* MCS Revocation Propagation Model - v0.2
\*
\* TLA+ specification of the five revocation effects declared in the
\* MCS paper §6.1 and formalised in LTL in Appendix F.
\*
\* Not yet accompanied by a published TLC transcript at v0.2. Model-checking with TLC requires a
\* bounded finite-state instantiation of Delegations, WorkUnits, and
\* Artefacts; the bound choices are documented in the README.
\*
\* Paper reference: §6.1 (five effects), Appendix F (LTL formulas).
\*
\* © 2026 J. Janssen Apparens, CC BY 4.0
\*

EXTENDS Naturals, FiniteSets, Sequences, TLC

CONSTANTS
  Delegations,      \* finite set of possible delegations
  WorkUnits,        \* finite set of possible work units
  Artefacts,        \* finite set of possible artefacts
  InitialSubDelegation, \* function Delegations -> SUBSET Delegations
  InitialWorkAuth,  \* function WorkUnits -> Delegations
  InitialProducedUnder, \* function Artefacts -> Delegations
  StopsInflightSet \* subset of Delegations with StopsInflight opt-in

ASSUME
  /\ IsFiniteSet(Delegations)
  /\ IsFiniteSet(WorkUnits)
  /\ IsFiniteSet(Artefacts)
  /\ InitialSubDelegation \in [Delegations -> SUBSET Delegations]
  /\ InitialWorkAuth \in [WorkUnits -> Delegations]
  /\ InitialProducedUnder \in [Artefacts -> Delegations]
  /\ StopsInflightSet \subseteqq Delegations

VARIABLES
  revoked,          \* SUBSET Delegations currently revoked
  workState,        \* [WorkUnits -> {"queued","in_flight","completed","cancelled"}]
  staleArtefacts,  \* SUBSET Artefacts marked stale
  clock            \* step counter used by bounded-latency variants

vars == <<revoked, workState, staleArtefacts, clock>>

\*-----
\* Transitive closure of InitialSubDelegation
\*-----

\* Reflexive-transitive closure: SubStar(d1,d2) iff d1=d2 or d2 reachable from d1 via sub-delegation.
\* We compute it iteratively rather than recursively for TLC compatibility.

ReachableFrom(d) ==
  LET
    Step(S) == S \union UNION { InitialSubDelegation[x] : x \in S }
    Iter[n \in 0..Cardinality(Delegations)] ==
      IF n = 0 THEN {d}
      ELSE Step(Iter[n-1])
  IN Iter[Cardinality(Delegations)]

SubStar(d1, d2) == d2 \in ReachableFrom(d1)

EffectiveRevokedBy(S, d) ==
  \E a \in S : SubStar(a, d)

EffectiveStopInflightBy(S, d) ==
  \E a \in S : SubStar(a, d) /\ a \in StopsInflightSet

\*-----
\* Initial state
\*-----

Init ==
  /\ revoked = {}
  /\ workState = [w \in WorkUnits |-> "queued"]
  /\ staleArtefacts = {}
  /\ clock = 0

\*-----
\* Actions
\*-----

\* Revoke a delegation directly (external action).
\* CancelsQueuedWork: queued work under d becomes cancelled in the next state.
\* StopsInflightWork: only under opt-in; otherwise in-flight continues.
\* MarksArtefactsStale: artefacts produced under d become stale in the next state.
RevokeDelegation(d) ==
  /\ d \in Delegations
  /\ d \notin revoked
  /\ LET NewRevoked == revoked \union {d} IN
     /\ revoked' = NewRevoked
     /\ workState' = [w \in WorkUnits |->
        CASE EffectiveRevokedBy(NewRevoked, InitialWorkAuth[w]) /\ workState[w] = "queued"
        -> "cancelled"
        ] EffectiveRevokedBy(NewRevoked, InitialWorkAuth[w]) /\ workState[w] = "in_flight" /\
EffectiveStopInflightBy(NewRevoked, InitialWorkAuth[w])
```

```

        -> "cancelled"
    [] OTHER
        -> workState[w]]
    /\ staleArtefacts' = staleArtefacts \union
      { a \in Artefacts : EffectiveRevokedBy(NewRevoked, InitialProducedUnder[a]) }
    /\ clock' = clock + 1

/* Propagate revocation transitively to direct sub-delegations of already-revoked delegations.
/* Each propagation step revokes one sub-delegation; bounded-latency properties require
/* the propagation to reach all descendants within k steps.
PropagateRevocation ==
    /\ \E d1 \in revoked, d2 \in InitialSubDelegation[d1] :
        /\ d2 \notin revoked
        /\ revoked' = revoked \union {d2}
        /\ workState' = [w \in WorkUnits |->
            CASE InitialWorkAuth[w] = d2 /\ workState[w] = "queued"
                -> "cancelled"
            [] InitialWorkAuth[w] = d2 /\ workState[w] = "in_flight" /\ d2 \in StopsInflightSet
                -> "cancelled"
            [] OTHER
                -> workState[w]]
        /\ staleArtefacts' = staleArtefacts \union
          { a \in Artefacts : InitialProducedUnder[a] = d2 }
        /\ clock' = clock + 1

/* Work unit progresses from queued to in_flight to completed.
/* v0.2 blocks work when its authority is effectively revoked through an ancestor.
/* This fixes the v0.1 direct-check gap, where a sub-delegation could start work
/* during the propagation interval.
BeginWork(w) ==
    /\ workState[w] = "queued"
    /\ ~EffectiveRevokedBy(revoked, InitialWorkAuth[w])
    /\ workState' = [workState EXCEPT ![w] = "in_flight"]
    /\ clock' = clock + 1
    /\ UNCHANGED <<revoked, staleArtefacts>>

CompleteWork(w) ==
    /\ workState[w] = "in_flight"
    /\ workState' = [workState EXCEPT ![w] = "completed"]
    /\ clock' = clock + 1
    /\ UNCHANGED <<revoked, staleArtefacts>>

Next ==
    /\ \E d \in Delegations : RevokeDelegation(d)
    /\ PropagateRevocation
    /\ \E w \in WorkUnits : BeginWork(w)
    /\ \E w \in WorkUnits : CompleteWork(w)

Spec == Init /\ [] [Next]_vars /\ WF_vars(PropagateRevocation)

/*-----
/* Properties (correspond to Appendix F.3)
/*-----

/* F.3.1 Transitive revocation (liveness): if d1 revoked and d2 in Sub*(d1),
/* then d2 eventually revoked.
TransitiveRevocation ==
    \A d1, d2 \in Delegations :
        ( d1 \in revoked /\ SubStar(d1, d2) ) -> ( d2 \in revoked )

/* F.3.3 CancelsQueuedWork (safety): once any ancestor a of d is revoked, no work
/* unit under d remains queued. v0.5.1 rewrite: quantifies over the transitive
/* closure (\E a \in revoked : SubStar(a, d)) rather than over revoked directly,
/* so that a buggy implementation dropping transitive propagation is detectable.
CancelsQueuedWork ==
    [] ( \A d \in Delegations :
        ( \E a \in revoked : SubStar(a, d) ) =>
            \A w \in WorkUnits :
                InitialWorkAuth[w] = d => workState[w] # "queued" )

/* F.3.4 StopsInflightWork under opt-in (safety):
/* once any ancestor a of d is revoked AND d in StopsInflightSet, no work unit
/* under d is in_flight. v0.5.1 rewrite: transitive-closure quantification.
StopsInflightUnderOptIn ==
    [] ( \A d \in Delegations :
        ( \E a \in revoked : SubStar(a, d) ) =>
            d \in StopsInflightSet =>
                \A w \in WorkUnits :
                    InitialWorkAuth[w] = d => workState[w] # "in_flight" )

/* F.3.4 Default case (safety): if d revoked AND d NOT in StopsInflightSet,
/* in-flight work under d may continue until completion. Expressed negatively
/* as: the model does NOT force d's in-flight work to cancel immediately.
InflightContinuesByDefault ==
    [] ( \A d \in revoked :
        d \notin StopsInflightSet =>
            \A w \in WorkUnits :
                ( InitialWorkAuth[w] = d /\ workState[w] = "in_flight" )
                => workState[w] \in {"in_flight", "completed"} )

/* F.3.5 MarksArtifactsStale (liveness + persistence):
/* every artefact produced under a delegation reachable from a revoked ancestor
/* eventually becomes stale, and once stale, remains stale.
/* v0.5.1 rewrite: transitive-closure quantification on the producing delegation.
ArtefactsEventuallyStale ==
    \A a \in Artefacts :
        ( \E anc \in Delegations :
            anc \in revoked /\ SubStar(anc, InitialProducedUnder[a]) )
        -> ( a \in staleArtefacts )

/* StalenessPersists is already direct on artefacts; transitive closure is
/* not relevant here because once an artefact is in staleArtefacts the property
/* is over the artefact itself, not its producing delegation.

```

```

StalenessPersists ==
  [] ( \A a \in Artefacts :
    a \in staleArtefacts => [] ( a \in staleArtefacts ) )

/*-----
/* Invariant that TLC can check in a single-step scan
/*-----

TypeOK ==
  /\ revoked \subseteqq Delegations
  /\ workState \in [WorkUnits -> {"queued","in_flight","completed","cancelled"}]
  /\ staleArtefacts \subseteqq Artefacts
  /\ clock \in Nat

/*-----
/* Notes on verification
/*-----
/*
/* At v0.1 this module is a specification, not a verified artefact.
/* A concrete model-checking run requires a finite instantiation.
/* Suggested smallest instantiation for sanity-check:
/*
/*   Delegations ← { d1, d2, d3 }
/*   WorkUnits ← { w1, w2 }
/*   Artefacts ← { a1 }
/*   InitialSubDelegation ← d1 ↦ {d2}, d2 ↦ {d3}, d3 ↦ {}
/*   InitialWorkAuth ← w1 ↦ d1, w2 ↦ d3
/*   InitialProducedUnder ← a1 ↦ d2
/*   StopsInflightSet ← {} or {d1}
/*
/* Expected outcomes at that scale:
/*   - TypeOK always holds
/*   - CancelsQueuedWork holds
/*   - StopsInflightUnderOptIn holds
/*   - TransitiveRevocation holds under WF_vars(PropagateRevocation)
/*   - ArtefactsEventuallyStale holds
/*   - StalenessPersists holds
/*
/* Publication of TLC verification transcripts remains required before any proof claim. This v0.2 file
fixes the direct-check gap but is still a specification until TLC output is published of
/* this artefact, independent of the paper version.

====

```

## Appendix K: TLC config

---

Source file: `tla_model_checking/mcs_revocation_model_v0_2.cfg`

```
SPECIFICATION Spec
INVARIANT TypeOK
PROPERTY TransitiveRevocation
PROPERTY CancelsQueuedWork
PROPERTY StopsInflightUnderOptIn
PROPERTY ArtefactsEventuallyStale
PROPERTY StalenessPersists

/* This cfg requires a companion constants instantiation or Toolbox model.
/* It is included to make the intended TLC checks explicit, not as a transcript.
```

## Appendix L: TLC status v0.5

---

Source file: `tla_model_checking/TLC_STATUS_v0.5.md`

## TLA+ / TLC status v0.5

---

Status: TLC model checking was not executed in this package.

Reason: the current execution environment does not provide `tla2tools.jar` / TLC. A TLC configuration skeleton is included so that Claude, Gemini, or a human reviewer can run the model externally.

Files: - `mcs_revocation_model_v0.2.tla` - `mcs_revocation_model_v0.2.cfg`

Claims allowed: - The model is a formal sketch of intended revocation semantics. - The v0.2 model includes `EffectiveRevokedBy` to block new work under direct or transitively revoked authority. - The model defines safety properties to be checked.

Claims not allowed until TLC output is attached: - "TLA+ proves revocation propagation." - "The temporal model has been model-checked." - "All revocation cases are safe."

Suggested command, assuming TLC is installed:

```
java -cp tla2tools.jar tlc2.TLC -config mcs_revocation_model_v0.2.cfg mcs_revocation_model_v0.2.tla
```

Required evidence before peer-review: 1. TLC command used. 2. TLC version. 3. Full TLC output transcript. 4. State-space size. 5. Whether invariants passed. 6. One negative-control model or old v0.1 counterexample showing the former transitive blocking defect.