

A Supervisory-Evidence Ontology for Agentic AI under EU Law

Jeroen Janssen, Apparens

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	Skan
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	<code>gov:Agent</code> with subtypes <code>SoftwareAgent</code> , <code>HumanAgent</code> , <code>OrganizationalAgent</code>	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	<code>gov:LegalRole</code> with eight AI Act and GDPR subtypes	Role-typed attribution; time-indexed via <code>gov:RoleAssignment</code>	GDPR Arts 4(7)-(8), 26, 28; AI Act Art. 3(3)-(8)	DPV present for GDPR roles; AI Act roles absent everywhere
3	<code>gov:Activity</code>	Agent action; inherits <code>startedAtTime</code> , <code>endedAtTime</code> , <code>wasAssociatedWith</code>	AI Act Art. 12	PROV-O present
4	<code>gov:ToolInvocation</code>	Subclass of <code>Activity</code> ; properties <code>hasTool</code> , <code>hasInvocationInput</code> , <code>hasInvocationOutput</code>	AI Act Art. 13	Agent Spec partial
5	<code>gov:MemoryEvent</code> with subclasses <code>MemoryRead</code> , <code>MemoryWrite</code>	Provenance of model/agent memory access	AI Act Art. 12; Art. 72	Atlan and Skan present; others partial
6	<code>dpv:DataSubject</code> (reused)	Data subject reference	GDPR Art. 4(1), 15, 22	DPV and GDPRov present; agent ontologies absent
7	<code>dpv:PersonalData</code> (reused)	Personal-data category hierarchy including Art. 9 special categories	GDPR Art. 9, 10, 30(1)(c)	DPV present; others absent
8	<code>dpv:Purpose</code> (reused)	Purpose of processing	GDPR Art. 5(1)(b), 6(4), 30	DPV present
9	<code>dpv:LegalBasis</code> (reused)	Lawful basis for processing	GDPR Art. 6, 9	DPV present
10	<code>dpv:Consent</code> plus <code>gdprov consent lifecycle</code> (reused)	Consent record with lifecycle	GDPR Art. 6(1)(a), 7	DPV and GDPRov present
11	<code>gov:Decision</code> with three subtypes (<code>SolelyAutomatedDecision</code> , <code>AssistedDecisionWithMeaningfulHumanIntervention</code> , <code>NonAutomatedDecision</code>), plus <code>gov:Score</code> and <code>gov:drawsStronglyOn</code> predicate	Three-way Article 22 taxonomy plus upstream attribution	GDPR Art. 22(1); SCHUFA para 73; Dun and Bradstreet para 40	Absent everywhere
12	<code>gov:HumanIntervention</code>	Six-property meaningful-intervention schema: <code>hasAuthorityToDeviate</code> , <code>hasUnderstandingOfLogic</code> , <code>consideredAllRelevantData</code> , <code>hasMarginOfDiscretion</code> , <code>interventionDurationSeconds</code> , <code>authorityActuallyExercised</code> (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-statute 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_v0_3.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 ≥ 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.

```
V i. 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 $F(\text{operational_disruption} \wedge \text{essential_services})$ 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 \times NIS2 incident window, F12 memory write \times NIS2 Art. 23(4) reporting cadence, F13 policy version change \times 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, μ -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 ≥ 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 ≤ €50 million or annual balance sheet total ≤ €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 ≤ €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 "P1Y"^^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 Δt 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(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 (ϵ , M); Profile B uses a qualitative instantiation aligned with national-court practice.

Profile A quantitative. drawsStronglyOn(d, score, ϵ , M) holds iff the empirical flip rate of (score, d) under counterfactual perturbation magnitude ϵ exceeds threshold M. Computable from PROV-O DAGs via the procedure: for each score s used by decision d, perturb s by ϵ across its empirical distribution; re-run the decision pipeline; measure the proportion of perturbation replicas in which d flips outcome. Parameters: ϵ = 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. Indicium 1: the score uses a specialised scoring model (Wiesbaden VG January 2026 treated this as itself evidence of strong reliance). Indicium 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). Indicium 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 (ϵ = 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 κ 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/GDPProv, 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;	Solely	Judgment paras 40, 59-60

			"lacked any manual oversight".		
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; Krakow review "not much more than a purely symbolic act".	Solely (REVERSED first instance Assisted)	ECLI:NL:GHAMS:2023:793
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	Out-of-Art.22-Scope	FPF Cases 10-13

			images; decided on Art. 6/9.		
B19	CNIL / Datainspecties 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 Belastingdienst Toeslagen	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 Belastingdienst toezichtarrangement	Aug 2025	AP requested plan-of-action on selectie-instrumenten.	Pending	AP letter 8 Aug 2025
B22	District Court The Hague SyRI	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 Mevaluate appeal	25 May 2021	Confirmed Garante.	Solely	Cassazione 25 May 2021
B24	CNIL Clearview fine	Oct 2022	€20M fine; decided on Art. 6/9.	Out-of-Art.22-Scope	CNIL Oct 2022
B25	Hamburg DPA SCHUFA (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 < 100ms 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.

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

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).

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

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.

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

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