EVIDE External Domain Contribution

HR & AI Governance — Schema Development v1.2 → v1.8

EXTERNAL DOMAIN CONTRIBUTOR	FRAMEWORK AUTHOR
Saly Man *AI Governance Architect* EU AI Act Compliance for HR & Recruitment AI Systems Defensible AI Decisions, Risk Registers, Human Oversight & Audit Logs linkedin.com/in/saly-ai-governance/	**Dott. Emanuel Celano** *Protocol </AI> Founder | EVIDE* Digital Evidence & AI Governance Informatica in Azienda — Bologna, Italy CertifyWebContent.com | DAPI-Certification.com linkedin.com/in/emanuelcelano/
Period: March – April 2026	**Framework:** EVIDE / </AI> Protocol

Context

This document records the external domain contribution of Saly Man to the development of the EVIDE framework (External Verifiable Integrity of Decision Events) and the </AI> Protocol, produced by Informatica in Azienda under the direction of Dott. Emanuel Celano.

The contribution was developed through a structured technical exchange focused on the operational application of EVIDE to HR and AI governance workflows, including AI-assisted candidate screening, human override scenarios, escalation paths, and classification governance under the EU AI Act.

Schema Contributions — v1.2 to v1.8

v1.2	`intervention.rationale_type`	Classification layer for intervention types, introduced following feedback that rationale was consistently present but narrative and non-comparable across cases.
v1.3	`intervention.taxonomy_version`	Anchoring the classification context at the moment of deposit, enabling classification replay under audit even as taxonomies evolve over time.
v1.4	`fedis_requested + API ingest`	FEDIS integration path and programmatic API ingestion, enabling scalable and automated evidentiary deposit flows.
v1.5	`authority_verification_status`	Server-computed response attribute clarifying whether authority verification was claimed or declared at intake, improving transparency at the boundary.
v1.6	`intervention.classification_status`	Structural state of classification at deposit (stable / provisional / contested), exposing divergence between reviewers without forcing artificial convergence.
v1.7	`intervention.classification_context`	Taxonomy reference, threshold reference, and threshold status — exposing whether admissibility had a defined structure upstream, introduced following the observation that most governance gaps emerge at the absence of a defined decision condition.
v1.8	`threshold_authority (attribution_status)`	Threshold ownership attribution layer — exposing whether a threshold had a single attributable source of authority, was fragmented across functions, implicit in practice, or unknown at closure. Introduced following the observation that authority fragmentation creates a different evidentiary condition than simple absence of a threshold.

Key Alignment Points

— **Intervention traceability vs. decision accountability**

Clarified the structural separation between tracing the path of human intervention (internal traceability) and anchoring accountability at the decision outcome (evidentiary layer). This distinction became foundational to the EVIDE architecture.

— **Anchoring threshold as an operational decision**

Identified that the boundary at which a decision is considered 'sufficiently defined to anchor' is not a technical constraint but an operational governance decision with durable consequences.

— **Taxonomy drift and inter-reviewer consistency**

Introduced the distinction between classification drift over time (addressed by taxonomy_version) and divergence at entry (addressed by classification_status and lightweight guidance at the point of classification).

— **Authority fragmentation vs. authority absence**

Identified that threshold ownership can be fragmented across policy, operational, and model governance functions without being absent, creating a different evidentiary condition from 'not_defined' — directly leading to the threshold_authority field in v1.8.

— **Authority incoherence at the closure point**

Introduced the distinction between fragmented authority and conflicting authority — where competing conditions cannot all be satisfied simultaneously, creating ambiguity in responsibility attribution at the decision closure boundary. Identified as a candidate for explicit modeling in a future EVIDE iteration.

Public Reference

This contribution is referenced in the External Signals section of the EVIDE public documentation under the category 'External Expert Signals — HR Governance Scenario.'

Public documentation: app.certifywebcontent.com/signals

EVIDE JSON Schema: app.certifywebcontent.com/json