



ARTISANAL INTELLIGENCE

Research Companion

v0.6 Release Candidate

English edition

Website publication edition · 17 August 2026

PUBLIC PILOT EDITION · NOT CERTIFICATION

Research status. This companion documents the design lineage, external relationships, limitations and validation agenda of the v0.6 Release Candidate. It is explanatory and research-oriented. It does not assert that the framework has already achieved scientific validation, legal conformity, certification readiness or universal novelty.

1. Origin and research question

Artisanal Intelligence began with a practical concern: rapid adoption of AI and automation can create useful outputs while weakening an organisation's ability to know where its knowledge came from, why a decision was made and how the work can continue when people, systems or suppliers change.

The project therefore moved away from governing AI as the primary object. Its research question became:

How can an organisation demonstrate that crucial knowledge retains sufficient integrity, traceable provenance, clear authority, practical understandability and recoverable capability across technological, human, supplier and organisational change?

AI remains important, but as one Knowledge Risk Driver alongside outsourcing, retirement, reorganisation, migration, poor archiving, data-quality decline and concentrated dependency.

2. Two framing concepts

2.1 Knowledge Pollution

Knowledge Pollution describes the gradual degradation of organisational knowledge when unsupported, stale, decontextualised or generated content is repeatedly reused until its origin, limitations and authority are obscured. The concept does not claim that generated content is inherently poor. The risk arises when transformations and reuse outpace traceability, challenge and correction.

Observable patterns include:

- generated summaries becoming de facto sources while primary material disappears;
- several dashboards reproducing the same upstream data defect;
- procedures retaining steps while losing the professional reasoning behind exceptions;
- copied text acquiring authority merely through repetition;
- a migrated repository preserving files while losing metadata, links, rights or interpretive context.

2.2 Knowledge Resilience

Knowledge Resilience is the ability to preserve, access, understand, transfer, reconstruct and responsibly apply the knowledge capability required for a function when conditions change. Storage is relevant but insufficient. A resilient organisation can demonstrate operational recovery, not merely point to documents or escrow.

3. Design evolution and contribution discipline

The classifications below are deliberately conservative.

Design element	Evolution	Classification at v0.6 RC
Knowledge-assurance focus	Shifts the primary question from managing an AI system to demonstrating the condition of organisational knowledge.	New synthesis.
Knowledge Pollution	Names degradation through unsupported or circular reuse, provenance loss and authority laundering.	Proposed original framing; requires empirical discrimination from adjacent concepts.
Knowledge Resilience	Connects preservation and continuity to operationally recoverable knowledge capability.	New synthesis.
Three knowledge types	Source, Craft and Generated distinguish functional roles in a knowledge chain.	New synthesis; not claimed as an exhaustive ontology.
Five assurance properties	Integrity, Provenance, Authority, Understandability and Resilience form one evidence-bounded lens.	New synthesis.
Knowledge Evidence Architecture	Treats evidence as cross-cutting rather than a sixth property.	New synthesis.
Awareness Principle	Makes unknown or falsely assumed dependencies the primary assurance concern while tolerating bounded known uncertainty.	Proposed original contribution pending reproducibility testing.
Minimum Sufficient Evidence	Applies proportional claim-specific evidence instead of maximal documentation.	New synthesis drawing on established assurance logic.
Knowledge Continuity Reserve	Defines a verified package enabling reconstruction, transfer or continuity.	Proposed original mechanism pending live validation.
Knowledge Recovery Objective	Links a context-specific recovery time to an explicit knowledge-capability criterion.	Proposed original mechanism pending calibration.
Optional Assurance Modules	Keeps the core stable while allowing specialised, automated and standards-mapping methods at the edge.	Architectural synthesis.

The distinctive claim is not ownership of familiar knowledge-management, records, continuity or provenance concepts. It is the repeatable movement from **dependency and consequence** to an **evidence-bounded, property-specific, time-bounded assurance conclusion**.

4. Rejected or constrained directions

- **AI as governance object** - rejected because knowledge risks pre-exist and extend beyond AI.
- **A sixth assurance property** - rejected because evidence architecture supports all properties.

- **Exhaustive knowledge ontology** - rejected; the three types are role-relative functional categories.
- **Generic competence threshold** - rejected; competence is contextual evidence for Authority, Understandability and Resilience.
- **Escrow or documentation as proof of resilience** - rejected; availability without interpretability and capability is insufficient.
- **Maximal evidence or checklist assurance** - rejected in favour of Minimum Sufficient Evidence.
- **Composite precision score for conformance** - rejected because it can hide a material property weakness.
- **Mandatory job title** - rejected; accountable authority is functional and context-sensitive.
- **Mandatory Assurance Modules** - rejected; a material trigger must be addressed, but the mechanism remains proportionate and optional.

5. Relationship to adjacent disciplines

Discipline or framework family	What it contributes	Artisanal Intelligence emphasis
Knowledge management	Organisational knowledge strategy, culture, learning, sharing and management systems.	Evidence-bounded assurance of a defined knowledge-dependent object and intended use.
Records management and archives	Authenticity, capture, metadata, retention, preservation and long-term access.	Adds Craft and Generated Knowledge, operational understanding and recovery capability without replacing records rules.
Business continuity	Impact analysis, continuity strategies, response and recovery.	Focuses on the knowledge capability and context that must survive or be restored.
Information security and cloud governance	Confidentiality, integrity, availability, controls and supplier responsibility.	Examines whether secure or available information remains authoritative, understandable and operationally usable as knowledge.
AI governance	AI-specific risk, transparency, human oversight, lifecycle and legal requirements.	Treats AI as one driver and follows its effect on organisational knowledge. Other AI obligations remain applicable.
Data governance	Data quality, access, rights, stewardship and lifecycle.	Distinguishes data and records from the professional judgement and generated products that make a knowledge-dependent function work.
Provenance standards	Entities, activities, agents and traceable transformation.	Uses provenance as one property within a broader assurance conclusion.
Audit and assurance	Claim, criteria, evidence, professional judgement, independence and reporting discipline.	Adapts those disciplines to knowledge assurance while explicitly withholding external audit or certification claims at v0.6 RC.

6. External foundations and mapping snapshot

This mapping is informative. It records relationships; it does not assert equivalence, legal compliance or automatic transfer of certification. Snapshot date: **17 August 2026**.

Source	Relationship to Artisanal Intelligence	Official source
EU AI Act - Regulation (EU) 2024/1689, as amended	AI-system obligations, governance, transparency, human oversight and risk. The knowledge-assurance framework remains complementary.	Official source
Regulation (EU) 2026/1744 - Digital Omnibus on AI	Amends application dates and selected AI Act requirements; demonstrates why external mappings must be decoupled from the constitutional core.	Official source
Interoperable Europe Act - Regulation (EU) 2024/903	Public-sector interoperability includes organisational, semantic and technical dimensions relevant to knowledge transfer.	Official source
Data Governance Act - Regulation (EU) 2022/868	European data-governance and reuse context; not a substitute for knowledge assurance.	Official source
Data Act - Regulation (EU) 2023/2854	Access, use and switching context relevant to supplier dependency and portability.	Official source
Open Data Directive - Directive (EU) 2019/1024	Public-sector information reuse context.	Official source
European Digital Identity Framework - Regulation (EU) 2024/1183	Trust services, electronic archiving and identity context.	Official source
ISO 30401:2018 and amendments	Knowledge-management system requirements; a related but broader management-system discipline. A replacement draft was under development at the snapshot date.	Official source
ISO 15489-1:2016	Concepts and principles for records creation, capture and management.	Official source
ISO 23081-1:2017	Principles for records metadata and provenance.	Official source
ISO 22301:2019 and amendments	Business-continuity management-system requirements; informs resilience but does not define the knowledge capability itself.	Official source
ISO 14721:2025 OAIS	Long-term preservation and Designated Community understandability context.	Official source
ISO/IEC 42001:2023 and amendments	AI management systems; complementary AI-specific management discipline.	Official source
ISO/IEC 27017:2026	Cloud information-security controls and shared customer-provider responsibility; relevant to supplier/platform dependency.	Official source
NIST AI RMF 1.0 and GenAI Profile	Voluntary AI risk-management resources. NIST stated at the snapshot date that AI RMF 1.0 was being revised.	Official source
OECD Governing with Artificial Intelligence (2025)	Public-sector AI benefits, risk and implementation context.	Official source
W3C PROV	Interoperable provenance concepts for entities, activities and agents.	Official source

7. 2026 watch notes

7.1 EU AI Act application dates

The consolidated AI Act applies generally from 2 August 2026, but Regulation (EU) 2026/1744 changed important dates. At the snapshot date, Chapter III Sections 1-3 were scheduled to apply from **2 December 2027** for Annex III high-risk systems and **2 August 2028** for high-risk systems linked to Annex I product legislation. Providers and deployers of existing high-risk systems intended for public authorities have a separate transition to **2 August 2030** under the amended Article 111.

These moving dates are a mapping issue, not a constitutional trigger. Any legal application requires case-specific analysis of the consolidated text.

7.2 Standards lifecycle

- ISO 30401:2018 remained current with amendments while a replacement draft was under development.
- ISO 22301:2019 remained the published base standard, with a committee revision in development.
- ISO/IEC 27017:2026 was published in July 2026 and replaced the 2015 edition.
- NIST reported that AI RMF 1.0 was being revised.

The Research Companion should be versioned more frequently than the Constitution because external sources change.

8. Methodological limitations

- Independent external multi-assessor reproducibility testing is not complete.
- The Basic, Enhanced and Critical profiles are reasoned pilot profiles and require empirical calibration.
- The Awareness Principle may still produce inconsistent judgement between “known but incomplete” and “unknown or falsely assumed.”
- KRO and the Knowledge Continuity Reserve require live recovery and transfer evidence across public-service contexts.
- The Generated Knowledge boundary requires calibration in deterministic, analytical and human-AI co-authored cases.
- Minimum Sufficient Evidence may be interpreted too lightly unless claim-specific sufficiency and correlation are tested.
- The module interface improves disclosure but does not itself prove module quality.
- Administrative burden and incremental value over existing public-sector controls remain to be measured in live pilots.
- v0.6 RC establishes neither legal compliance nor an accreditation basis.

9. Validation and falsification agenda

The method should be tested against friendly, failure and adversarial cases.

Test	Core question	Concerning result
Independent assessor test	Do assessors find materially similar dependencies, profiles, property states and reliance outcomes?	Opposite reliance positions without a traceable ambiguity.
Hidden-dependency test	Does active discovery reveal Craft Knowledge, supplier transformations and exception paths absent from the initial description?	Repeated omission of seeded material dependencies.
Adverse-evidence test	Does late contradictory evidence reopen the provisional position?	Anchoring preserves the original conclusion despite material contradiction.
Evidence-cloning test	Are several derivatives from one flawed source recognised as correlated?	Multiplicity is mistaken for independent strength.
Recovery test	Can another competent party perform representative work within the KRO and capability criterion?	Documentary review passes while operational recovery fails.
Ceremonial-review test	Does human review actually challenge AI output?	A tick-box reviewer is treated as strong evidence without competence, access or challenge.
Scope-gaming test	Does the boundary rule prevent strategic exclusion or unlimited expansion?	Material supplier or downstream dependencies are excluded; or scope becomes unusable without affecting the conclusion.

Early validation should use simple compatibility measures before relying on aggregate statistics. A calibration session should distinguish assessor error, case ambiguity, template ambiguity and method ambiguity. Repair the method only when divergence is systematic or unresolved from the existing rule.

10. Research hypotheses for public-sector pilots

1. The three knowledge types reveal dependencies that a document-only review misses.
2. Property-specific conclusions prevent important weaknesses from being hidden by an overall score.
3. Explicit awareness states make false certainty more visible than conventional risk-register wording.
4. Minimum Sufficient Evidence reduces burden when existing controls are mapped at claim level.
5. A capability-based KRO produces better continuity decisions than “documents available.”
6. Active dependency discovery materially improves assessment quality.
7. A One-Page Assessment can support management decisions without becoming an evidence ceiling.
8. Optional modules improve depth or efficiency only when material triggers and limitations are recorded.

11. Evidence Gap Register

Gap	Priority	Required work	v0.6 RC status
EG-001 Reproducibility	Critical	Blind multi-assessor tests on the same case set, followed by calibration and retest.	Open.
EG-002 Public-sector pilot	High	Complete at least one live municipality or agency pilot and measure decision value and burden.	Open.
EG-003 Consequence calibration	High	Test the three profiles against cases involving rights, continuity, legal effect and reversibility.	Open.
EG-004 Recovery calibration	High	Collect capability criteria, KROs and operational results by service context.	Open.
EG-005 Generated Knowledge boundary	Medium	Test deterministic analytics, rules engines, derived outputs and human-AI artefacts.	Open.
EG-006 Evidence equivalence	Medium	Build a claim-level library across records, interviews, observation, logs and capability tests.	Open.
EG-007 Module quality	Medium	Validate module disclosure, compatibility, change control, conflicts and revocation.	Open.
EG-008 Sector profiles	Low	Develop examples for permitting, grants, inspections, taxation, social services and crisis response.	Deferred to profiles.

12. Publication and change discipline

The constitutional core should not be reopened because an external standard changes or because a new sector example appears. External mappings, guidance, module specifications and validation protocols may evolve at the edge.

Claims of originality should remain provisional until compared systematically with adjacent academic and professional literature. Claims of usefulness, measurability and reproducibility should follow evidence from real pilots, not document maturity.

The next challenge of artificial intelligence is not only to control the technology, but to preserve the human knowledge on which reliable organisations are built.