



ARTISANAL INTELLIGENCE

Artisanal Intelligence Constitution

v0.6 Release Candidate

English edition

Website publication edition · 17 August 2026

PUBLIC PILOT EDITION · NOT CERTIFICATION

Publication status. v0.6 is a public Release Candidate for pilots and structured external review. It is not an external conformity standard, an accredited certification scheme, an audit opinion or a legal-compliance determination. The constitutional baseline was frozen on 15 August 2026; empirical validation remains open.

1. Purpose

Artisanal Intelligence is a knowledge-assurance framework. It helps an organisation demonstrate that crucial organisational knowledge retains sufficient **Integrity**, has traceable **Provenance**, is governed through clear **Authority**, remains **Understandable**, and is **Resilient** when people, technology, suppliers or organisational arrangements change.

The central question is:

How can an organisation demonstrate that its crucial knowledge remains trustworthy, traceable, authorised, understandable and recoverable when AI, outsourcing, staff changes, system migrations or other Knowledge Risk Drivers occur?

The governance object is organisational knowledge, not AI. AI is one Knowledge Risk Driver among many.

2. Scope and applicability

An assessment defines a **Knowledge Assurance Object**: a bounded knowledge-dependent function, process, service, decision domain, dataset, system-supported activity, organisational unit or collaboration. The record states its intended use, boundaries, material dependencies, lifecycle stages, exclusions and assessment period.

The framework is designed first for public administrations and other knowledge-intensive organisations, but it may be used elsewhere. It complements and does not replace applicable law, professional duties, records management, knowledge management, business continuity, information security, data governance, AI governance, safety or sector-specific controls.

Boundary and stopping rule. Include dependencies that could reasonably change the consequence profile, evidence sufficiency, assurance conclusion or required action for the stated intended use. Discovery may stop when additional knowledge elements would not reasonably change any of those outcomes. Record the stopping rationale and every material exclusion.

3. Normative language within v0.6

The terms below control claims made under the v0.6 pilot method. They do not turn this Release Candidate into an accredited external standard.

Term	Meaning
shall	Framework requirement for an assessed-conformance claim under v0.6.
should	Recommendation; departure requires context-sensitive judgement.
may	Permission.
can	Possibility or capability.

4. Constitutional principles

4.1 Knowledge First

Assurance shall start from the knowledge required for trustworthy action, service or decision. Within Artisanal Intelligence, technology is examined through its effect on that knowledge; independent legal, safety, security and ethical obligations remain fully applicable.

4.2 Awareness Principle

Proportionate awareness of material knowledge dependencies, uncertainties, limitations, losses and gaps is the baseline. Awareness is evidence-sensitive:

1. **Known and evidenced**
2. **Known but incompletely evidenced**
3. **Unknown or falsely assumed**

State 3 is the primary assurance concern. Retrospective discovery of a gap is not automatically failure when the limitation is recognised, assessed and managed proportionately. False certainty is more serious than acknowledged uncertainty.

4.3 Trust and challengeability

Trust is the starting point; challengeability is the safeguard. Self-declaration may contribute at low consequence. Evidence and genuine challenge shall become stronger as consequence, uncertainty and dependency increase.

4.4 Stable core

The three knowledge types and five assurance properties are the protected constitutional core. Knowledge Evidence Architecture is cross-cutting and is not a sixth property. No silent architectural change is permitted.

4.5 Consequence-based proportionality

Assurance depth, evidence strength, challenge and independence shall reflect reasonably foreseeable consequence, uncertainty, methodological limitation and dependency. The target is **Minimum Sufficient Evidence**, not maximal documentation.

4.6 Evidence equivalence

Different evidence may support the same assurance claim when relevance, credibility, scope, freshness, challengeability, proportional independence and limitations provide comparable support. Equivalence is claim-specific and shall not lower Minimum Sufficient Evidence. Nominally different evidence that depends on the same flawed upstream source does not gain strength through multiplicity.

4.7 No compensating average

All five properties remain separately visible. A material weakness in an essential property shall not disappear into a composite score or be cancelled by strength elsewhere.

4.8 Neutrality and regulatory decoupling

The framework is technology- and vendor-neutral. External legal, regulatory and standards changes normally update mappings and guidance rather than the constitutional core.

4.9 Six-Month Stability

The constitutional explanation should remain stable enough to be presented substantially unchanged for at least six months. This is a design-discipline rule, not an automatic validity period for an assessment.

4.10 Open core

Essential framework meaning and assessment logic shall not depend exclusively on proprietary professional services. Open framework material remains distinct from implementation support, training, research partnerships and future independent-assurance services.

5. Knowledge architecture

The three types are functional categories, not an exhaustive ontology. Classification is relative to the assessed dependency and the role an item plays in a knowledge chain.

Knowledge type	Constitutional meaning	Typical assurance concern
Source Knowledge	Knowledge embodied in or directly derived from primary, authentic or formally recognised sources, including law, official decisions, authentic records, original datasets, approved procedures and adopted specifications.	Source status does not itself prove correctness, currency or applicability.
Craft Knowledge	Professional and contextual knowledge needed to interpret, apply, adapt or challenge other knowledge and to perform a function responsibly; including experience, exceptions, reasoning, judgement, institutional history and tacit practice.	Assurance does not require documenting everything. Observation, demonstrated performance, peer challenge, redundancy, transfer and recovery tests may provide evidence.
Generated Knowledge	A knowledge product created, inferred, synthesised, recommended or materially transformed through AI, automation, analytics, modelling or another computational process and capable of informing judgement, decision, explanation, action or continuity.	Purely mechanical deterministic transformation need not be classified as Generated Knowledge when classification adds no assurance value.

A generated or human-produced artefact may later serve as Source Knowledge for a downstream dependency. That role change never resets inherited Provenance, Integrity, Authority, uncertainty or known limitations. Mixed artefacts may contain all three types.

6. Assurance properties

Property	Requirement
Integrity	Knowledge is sufficiently reliable, complete for purpose, current where needed, and protected against uncontrolled degradation, unsupported reuse, contamination, unauthorised change or material loss of context.
Provenance	Origin, contributors, dates, versions, transformations, dependencies, validation and material AI or automation use can be traced far enough to challenge the assurance conclusion.
Authority	Mandate and responsibility to establish, validate, approve, use, change, supersede and challenge knowledge are clear.
Understandability	An appropriately competent person can understand content, context, assumptions, reasoning, exceptions, limitations, dependencies and intended use well enough to act responsibly.
Resilience	The organisation can preserve, access, restore, reconstruct, transfer and apply critical knowledge during change or disruption.

7. Knowledge Risk Drivers

Knowledge Risk Drivers include AI and automation; staff turnover and retirement; outsourcing; supplier dependency; system migration; reorganisation; weak archiving; obsolete procedures; crises; access loss; fragmented authority; data-quality deterioration; technology obsolescence; and concentration of knowledge in one person, system or supplier.

A driver is not automatically a deficiency. The organisation considers the affected knowledge, properties, consequences, safeguards, evidence and residual uncertainty.

8. Knowledge-assurance lifecycle

1. Identify and bound the Knowledge Assurance Object.
2. Discover, acquire or create and classify relevant knowledge.
3. Interpret, validate and authorise knowledge.
4. Use and, where material, transform knowledge.
5. Retain, preserve and keep necessary context.
6. Transfer or migrate knowledge when dependencies change.
7. Restore or reconstruct knowledge capability when disrupted.
8. Withdraw or dispose only under applicable legal, archival, privacy and continuity rules.

9. Governance and accountability

The organisation remains accountable for knowledge used in its functions, including outsourced or automated work. An **Accountable Knowledge Authority** shall be identifiable at a proportionate level. The role is functional; no universal job title is prescribed.

The role shall have sufficient practical access and mandate to obtain evidence, challenge conclusions and escalate unresolved limitations. Knowledge-assurance accountability may remain distinct from statutory decision authority and business or policy risk acceptance; the relationship must be explicit.

10. Knowledge Evidence Architecture

Evidence connects assurance claims to observable support and visible limitations. It may include records, source references, metadata, logs, interviews, observation, samples, demonstrations, professional review, validation records, contracts, configurations and reconstruction, migration or transfer tests.

Evidence shall be relevant, credible, attributable, retrievable, sufficiently current, proportionate and challengeable. The record states principal evidence, claim linkage, material limitations, conflicts, common upstream dependencies, unresolved uncertainty and why further evidence was or was not required.

11. Consequence-based assurance

Before selecting assurance depth, the organisation evaluates reasonably foreseeable consequences, including citizen impact, legal effect, fundamental rights, safety, financial effects, public-service continuity, societal scale, reversibility, duration, detectability, automated-decision dependency and concentration of dependency.

v0.6 pilots may use **Basic**, **Enhanced** and **Critical** as provisional assessment-intensity profiles. They are not maturity grades, public labels or mandatory constitutional conformance levels. Their calibration remains an open validation task.

12. Resilience mechanisms

12.1 Capability acceptance criterion

For a material recovery dependency, define the minimum combination of knowledge, understanding, access, rights and competent capability required to perform the function without unacceptable consequence.

12.2 Knowledge Recovery Objective

The **Knowledge Recovery Objective (KRO)** is the context-specific period within which that capability acceptance criterion must be met. No universal recovery time is prescribed. The relevant boundary may be immediate, within a workday, before the next decision cycle, before a statutory deadline or another unacceptable-consequence point.

12.3 Knowledge Continuity Reserve

The **Knowledge Continuity Reserve (KCR)** is the minimum verified package of knowledge, artefacts, access rights and capabilities through which a critical function can be reconstructed, transferred or continued without unacceptable dependence on one supplier, technology or individual. Escrow or documentary availability alone is insufficient. Operational reconstruction, migration or transfer testing is stronger evidence where consequence warrants it.

13. Assurance Modules: stable core, extensible edge

Assurance Modules are optional extensions that may deepen, automate or specialise an assurance question. A valid lightweight assessment remains possible without a module. The existence of a module does not create a conformance requirement. A module-free assessment may nevertheless conclude that assurance is insufficient or that uncertainty remains unacceptable.

When a material trigger is present, the assessment records whether a relevant module was used, replaced by equivalent evidence, or not used, and why. A trigger is material when ignoring it could reasonably change consequence, evidence sufficiency, the assurance conclusion or required action.

Every reusable module declares:

- Purpose
- Trigger
- Assurance properties addressed
- Inputs
- Method
- Evidence or output
- Limitations

Modules may use human, automated, algorithmic or AI-assisted methods. Method type alone does not determine evidential strength. Accountability for the assurance conclusion remains with the responsible organisation or assessor. Additional independent verification is required only when consequence, uncertainty, methodological limitation or material dependence makes genuine separated challenge proportionate.

Modules may be core-maintained, reference or standards-mapping, or third-party. Category does not provide automatic evidential strength. Reusable modules should disclose compatible core versions, validation status, method and version identity, material dependencies and change history.

14. Assurance conclusions and claims

Conclusions address each property separately and disclose material deficiencies, evidence limitations, residual uncertainty, exclusions and conditions of validity. Formal v0.6 outcome vocabulary is controlled by the Assurance Manual; the Constitution does not create an external label.

Claim	Meaning	Restriction
Framework application	Selected concepts, methods or templates were used.	Not a conformance claim.
Assessed conformance	A defined Knowledge Assurance Object was assessed under the stated v0.6 method and profile.	Scope, intended use, period, assessor relationship, exclusions, limitations and validity conditions shall be disclosed.
Certification	Third-party claim under a formally governed scheme.	Not available under v0.6 RC. No certification or accreditation claim is authorised.

15. External requirements

Artisanal Intelligence is designed to complement law and established standards. External mappings are informative unless a stated assessment profile expressly incorporates them. A mapping or existing certificate may support claim-specific evidence; it does not automatically transfer compliance, certification or assurance.

16. Constitutional change control

A proposed change shall identify the problem, alternatives, evidence, affected clauses and artefacts, transition impact, bilingual impact and its Usefulness, Measurability and Reproducibility result. Changes to the three knowledge types, five properties or accepted constitutional principles require explicit approval by the initiative owner.

The core may be reopened only when validation demonstrates a genuine structural failure. Editorial clarification, evidence guidance, mappings, examples and module development remain possible without reopening the core.

Annex A. Protected boundary and open validation

Protected constitutional boundary: Organisational Knowledge as governance object; three knowledge types; five assurance properties; cross-cutting Knowledge Evidence Architecture; Knowledge First; Awareness Principle; trust with challengeability; proportionality; Minimum Sufficient Evidence; no compensating average; Regulatory Decoupling; Six-Month Stability; open core; controlled evolution.

Open validation gates: independent assessor reproducibility; live public-sector usefulness and burden measurement; consequence-profile calibration; recovery and KRO/KCR calibration; negative-case and gaming-resistance testing; Generated Knowledge boundary cases; and evidence-equivalence calibration.

Annex B. Version history

Version	Meaning
v0.1-v0.4	Exploration, concept expansion and architectural consolidation.
v0.5	Stable core: governance object, three knowledge types, five assurance properties and evidence architecture.
v0.6 RC	Operationalisation for pilots: awareness states, consequence profiles, evidence rules, resilience mechanisms, optional modules, assessment method and public-sector toolkit.