



ARTISANAL INTELLIGENCE

Assurance Manual

v0.6 Release Candidate

English edition

Website publication edition · 17 August 2026

PUBLIC PILOT EDITION · NOT CERTIFICATION

Method status. This manual operationalises the frozen v0.6 core for pilot assessment and external validation. It is authoritative for the internal v0.6 pilot method, but it is not an external conformity standard, accredited certification scheme, audit opinion or legal-compliance determination.

Make the assurance judgement explicit enough to test, challenge and reproduce - without turning the framework into a checklist.

1. Purpose and method boundary

The assessment determines whether a defined Knowledge Assurance Object demonstrates adequate awareness and assurance across **Integrity, Provenance, Authority, Understandability and Resilience** for its intended use and consequence context.

The method starts from organisational knowledge. AI, outsourcing, staff change, suppliers, migration and automation are relevant as Knowledge Risk Drivers. A procedurally valid assessment may still produce a negative outcome. Following the steps does not guarantee that the knowledge is sufficiently assured.

The method uses professional judgement, but bounds it through explicit scope, evidence, challenge, rationale, outcome and validity rules. There is no numeric averaging and no requirement to maximise documentation.

2. Assessment object, intended use and boundary

Each assessment defines a **Knowledge Assurance Object**: the task, decision, service, process, knowledge-dependent capability or other bounded object for which an assurance position is required.

Required element	Method rule
Object	Name the activity or decision being assessed; avoid vague organisation-wide labels unless the scope is genuinely organisation-wide.
Intended use	State what the organisation needs to do with the knowledge, for whom and under which operating conditions.
Boundary	Include dependencies that could reasonably change consequence, evidence sufficiency, conclusion or required action.
Material exclusions	Record exclusions that could reasonably affect any of those outcomes, with rationale.
Period	State the time and operating context to which the assessment applies.
Downstream use	Identify where an output becomes an input or Source Knowledge for another material dependency.

Stopping rule. Dependency discovery stops when additional knowledge elements would not reasonably change the consequence profile, evidence sufficiency, assurance conclusion or required action for the stated intended use. Record the stopping rationale.

3. Dependency discovery

The assessor must not assume that the initial process description contains all material knowledge. Discovery actively searches for hidden Craft Knowledge, exception paths, informal workarounds and supplier or system transformations.

1. Start from the task, decision or service and identify the outputs that must be produced responsibly.
2. Trace backwards to the Source, Craft and Generated Knowledge needed to produce, interpret or challenge those outputs.
3. Ask what experienced people do when the standard procedure does not fit.
4. Apply loss and change prompts: a person leaves; a source changes; a supplier exits; a system migrates; a model or rule changes; infrastructure fails; crisis conditions apply.
5. Check material transformations and hand-offs, including supplier processing, data pipelines, summarisation, human synthesis and automated generation.
6. Check downstream use. Reclassification changes functional role but never clears inherited provenance, integrity, authority, uncertainty or limitations.
7. Apply the stopping rule and record exclusions.

Craft Knowledge is not a documentation mandate. Observation, peer challenge, redundancy, demonstrated performance, transfer, simulation and recovery testing may provide stronger evidence than a newly written procedure.

4. Consequence model

For pilot comparability, v0.6 uses three provisional **assessment-intensity profiles**. They are not maturity levels, public labels or mandatory constitutional conformance levels.

Consider at least: citizen or user impact; legal effect or validity; fundamental rights; safety; financial or public-resource impact; public-service continuity; societal scale; reversibility; duration; detectability and correctability; automation dependency; and concentration of knowledge dependency.

Profile	Typical consequence context	Default evidence posture	Default challenge posture
Basic	Limited and normally reversible consequences; low dependency concentration; errors usually detectable and correctable.	Credible, current operational evidence may be light. Self-declaration can contribute when no material contradiction exists.	Targeted plausibility check; sample or observation when useful.

Profile	Typical consequence context	Default evidence posture	Default challenge posture
Enhanced	Material service, financial, legal or citizen consequences; notable dependency, automation or uncertainty.	Documented evidence for material claims; operational corroboration or triangulation when it reduces uncertainty.	Interview plus sample or observation; selected scenario, transfer or recovery test where relevant.
Critical	Potential serious rights, safety, legal-validity, major-continuity or large-scale consequences; high dependency concentration.	Strong, current and challengeable evidence; material controls or capabilities demonstrated in operation.	Mandatory challenge of material claims; independent verification and scenario, reconstruction or transfer testing when proportionate.

Start with the highest material consequence dimension. A mitigation may lower the profile only when the mitigation itself is evidenced and operational. Low perceived probability does not by itself neutralise a severe dependency-driven consequence.

5. Minimum Sufficient Evidence

Minimum Sufficient Evidence (MSE) is the least evidential basis still adequate to support the relevant assurance claim at the applicable consequence and uncertainty level. It is claim-specific, not a fixed document count.

Test	Question
Relevance	Does the evidence directly support the claim being made?
Credibility	Is the source or method sufficiently trustworthy for this claim?
Scope coverage	Does it cover the material part of the dependency and intended use?
Freshness	Is it current enough relative to the evidence cut-off and known changes?
Challengeability	Can a competent reviewer inspect, reproduce, observe or otherwise challenge it where proportionate?
Limitations	Are uncertainty, gaps and method limitations visible?
Correlation	Do apparently different items depend on the same upstream source, assumption, model or process?

Evidence is sufficient only when the material claim is supported for the intended use and remaining uncertainty is acceptable for the consequence profile. More evidence is not automatically stronger evidence. One direct operational test may be stronger than many derivative documents.

6. Evidence equivalence and correlation

Different evidence forms may be treated as equivalent only for the same assurance claim and only when relevance, credibility, scope, freshness, independence where proportionate and limitations provide comparable support.

- Equivalence never lowers Minimum Sufficient Evidence.
- Nominally different evidence derived from the same flawed source does not gain strength merely through multiplicity.
- A certificate, audit report or standards mapping is evidence only for the claims and period within its scope; assurance does not transfer automatically.

- When a proprietary or third-party method discloses too little about method, scope or limitations, that limitation may prevent sufficiency.
- Correlated evidence may still be useful, but must be treated as one dependency family rather than independent corroboration.

Example	Weak interpretation	Required interpretation
Five dashboards from one obsolete dataset	Five separate strong items.	One correlated evidence family with the same upstream defect.
ISO certificate plus internal policy	External certification proves the knowledge claim.	Test claim-level scope, period, evidence and gaps before reuse.
Manager statement plus copied procedure	Two confirmations.	Potentially one self-referential source unless practice is separately evidenced.

7. Five-property assessment criteria

Assess each property separately. The questions below are minimum prompts, not exhaustive checklists.

Property	Minimum assurance question	Typical red flags
Integrity	Are the knowledge and material sources sufficiently reliable, accurate, current and complete for the intended use?	Obsolete source; silent version change; completeness unknown; generated content accepted without validation.
Provenance	Can origin, transformations, material human or automated contributions and change history be traced far enough to challenge the conclusion?	Lineage breaks; provenance resets at migration; downstream Source Knowledge hides inherited limitations.
Authority	Are rights to create, approve, use, change, challenge and escalate knowledge clear, and do accountable roles have sufficient mandate?	Accountability without access or challenge rights; supplier or AI involvement displaces responsibility; exceptions have no owner.
Understandability	Can the people who must act understand meaning, context, limitations and relevant judgement well enough to use the knowledge responsibly?	Procedures describe steps but not decision logic; outputs lack meaning or limits; successor cannot handle exceptions.
Resilience	Can sufficient knowledge capability remain available or be restored in time when people, systems, suppliers or conditions change?	Documents exist but recovery fails; KRO has no capability criterion; rights or access disappear on supplier exit.

Strength in one property cannot automatically compensate for a material failure in another. If a property is essential to the intended use, the overall outcome cannot be stronger than its material limitation.

8. Awareness Principle in assessment

State	Meaning	Assessment treatment
1 - Known and evidenced	The dependency, uncertainty or limitation is recognised and supported by proportionate evidence.	Normally acceptable subject to evidence sufficiency and control.
2 - Known but incompletely evidenced	The issue is recognised, but evidence is incomplete, stale or weak.	May be acceptable at lower consequence when bounded; requires explicit limitation and stronger challenge as consequence rises.
3 - Unknown or falsely assumed	A material issue is not recognised or is treated as resolved without reasonable basis.	Primary assurance concern. At Critical profile, unresolved material State 3 normally prevents a positive outcome.

9. One-Page Assessment and Artisanal Case File

The One-Page Assessment is the visible conclusion layer. Detailed evidence and traceability may sit in the linked Artisanal Case File, provided the summary points to them unambiguously.

Visible core summary	Linked minimum traceability metadata
Knowledge dependency	Scope and intended use
Consequence profile	Assessor and accountable roles
Assurance outcome	Assessment date, version and evidence cut-off
Key strengths	Evidence references and Case File IDs
Key uncertainties	Planned review and change triggers
Actions	Module decision; residual uncertainty and limitations

“One page” is a communication constraint, not an evidence ceiling. Another competent assessor must be able to reconstruct the basis of the conclusion through the linked record.

10. Controlled assurance outcome model

The four-state model standardises v0.6 pilot wording. It may be used per property and for the overall outcome. It is not an external certification label.

State	Meaning	Reliance rule
Sufficiently assured	Evidence and residual uncertainty are adequate for the stated intended use and profile.	No material essential-property deficiency remains.
Sufficiently assured with limitations	Evidence is adequate for the bounded intended use, but explicit limitations, monitoring or planned action remain.	Limitations are bounded, owned and compatible with the consequence profile.
Insufficiently assured	A material deficiency or unacceptable uncertainty prevents responsible reliance.	Corrective action and reassessment are required before a positive reliance claim.
Not assessable	Evidence is too incomplete, contradictory or inaccessible to reach a responsible conclusion.	Not equivalent to sufficient; unresolved uncertainty constrains the decision.

No numeric averaging is permitted. An essential property that is Insufficiently assured or Not assessable normally caps the overall outcome accordingly.

11. Challenge and independent verification

Challenge tests whether the assessment survives a competent attempt to disconfirm it.

Technique	Best use
Interview	Reveal Craft Knowledge, awareness, authority, exceptions and implicit decision logic.

Technique	Best use
Observation	Verify that described practices occur in operation.
Sampling	Test representative records or cases without maximal review.
Scenario test	Test staff loss, supplier exit, outage, migration, source change or AI failure.
Reconstruction or transfer test	Have an alternate competent person or team re-establish the required knowledge capability.
Challenge review	A second person or function actively tests material claims, assumptions and limitations.

Independent verification is required only when consequence, uncertainty, methodological limitations or material dependence make it proportionate. “Independent” means sufficiently separated from the method, output or interest being verified to provide genuine challenge. No universal organisational separation is imposed.

12. Assurance Modules and material triggers

Modules are optional extensions. They add depth, automation or specialisation; they do not replace the core assessment and do not create requirements merely by existing. A valid module-free assessment can still produce an insufficient or not-assessable outcome.

A trigger is **material** when ignoring it could reasonably change consequence, evidence sufficiency, the assurance conclusion or required action.

Module	Typical trigger	Minimum decision record
AI Handling	AI materially contributes to analysis, classification, summarisation, recommendation or execution.	Module or equivalent evidence used; review effectiveness; provenance and dependency limits; change trigger.
Provenance Verification	A decision depends strongly on lineage or multiple transformations, systems or parties.	Chain tested; breaks and limitations; correlation or common-source issues.
Knowledge Recovery	Key-person, supplier, migration or continuity dependency could impair ability to act.	KRO when used; capability criterion; test result; unresolved gaps.
Automation and Monitoring	Assurance can decay quickly or manual rechecking is burdensome.	What is monitored; thresholds; owner; response and revalidation rule.
Standards Mapping	Existing control or audit evidence may support the same assurance claim.	Source framework, version and scope; claim-level mapping; gaps; no automatic transfer of certification.

Every reusable module declares **Purpose, Trigger, Assurance properties addressed, Inputs, Method, Evidence or output, and Limitations**. It also identifies compatible core versions, validation status, method/version identity, material dependencies and change history.

Modules may use human, automated, algorithmic or AI-assisted methods. Method type does not itself determine evidential strength. Accountability for the conclusion remains with the organisation or assessor. Additional independent verification follows consequence, uncertainty, method limitation and material dependence - not the mere use of AI.

Module categories are **core-maintained**, **reference or standards-mapping**, and **third-party**. Category controls provenance and governance, not automatic evidential strength.

13. Resilience, KRO and Knowledge Continuity Reserve

Resilience is operational capability, not document possession.

Element	Required meaning
Capability acceptance criterion	Explicit minimum knowledge capability required for the activity to resume without unacceptable consequence.
Knowledge Recovery Objective (KRO)	Context-specific target time by which the capability criterion must be met.
Knowledge Continuity Reserve (KCR)	Knowledge, artefacts, access rights, people or skills, instructions and other means retained to enable restoration.
Recovery test	Scenario, reconstruction or transfer exercise that demonstrates capability rather than access to documents.

Do not impose universal hours or days. Set the KRO against the next unacceptable-consequence point, such as a statutory deadline, payroll run, incident-response window, permit cycle or essential-service threshold.

14. Assessment validity and revalidation

An assurance conclusion is time-bounded.

Minimum validity record	Rule
Assessment date and version	Identify the exact assessment position and method version.
Evidence cut-off	State the latest date or event included in the evidence base.
Planned review point	Set a period proportionate to consequence and rate of change; no universal interval.
Event triggers	Revalidate after material change in source, process, supplier, system, authority, model, dependency, incident or contradictory evidence.
Monitoring trigger	Where monitoring is used, define what change crosses the reassessment threshold.
Supersession	A later assessment clearly supersedes or narrows the earlier one; stale positive claims must not remain ambiguous.

15. Roles, accountability and mandate

Role	Minimum responsibility
Assessor	Applies the method, records rationale, identifies uncertainty and does not overstate evidence.
Accountable Knowledge Authority	Owns the knowledge-assurance position and has sufficient mandate to obtain evidence, challenge conclusions and escalate unresolved limitations.

Role	Minimum responsibility
Operational owner	Implements controls and actions and maintains the knowledge-dependent activity.
Risk or statutory decision authority	Accepts business, policy or statutory risk where applicable; may differ from the Accountable Knowledge Authority.
Independent challenger	Provides proportionate genuine challenge when required; independence is functional and case-specific.

A named role without practical access or escalation power is an accountability sink and does not satisfy the method.

16. Assessor calibration and reproducibility

Reproducibility means materially compatible conclusions from the same case, not identical prose or evidence sets.

Difference	Interpretation
Minor	Different wording or evidence mix, but the same material dependencies, profile, property states, overall outcome and actions.
Acceptable bounded divergence	One adjacent state or profile difference with the same facts and uncertainty explicitly described; rationale and calibration required.
Major divergence	Positive versus insufficient or not-assessable outcome; missed material dependency; opposite conclusion on an essential property; ignored trigger; consequence differs by more than one profile.
Critical miss	A material failure is missed and inappropriate reliance results, especially in a Critical case.

Calibration compares scope first, then consequence, evidence sufficiency, property conclusions and finally overall wording. Persistent ambiguity is recorded as a method-defect candidate.

17. Core evidence catalogue

The catalogue gives default evidence patterns, not a mandatory document list.

Evidence item	Properties	Minimum content	Verification	Typical limitation
Scope and boundary record	All	Object, intended use, dependencies, exclusions, period, stopping rationale.	Owner confirmation and challenge of exclusions.	Scope shaving may hide dependencies.
Source register	Integrity; Provenance	Authoritative source, version or date, owner, applicability.	Trace a sample to source.	A list does not prove currency.
Craft Knowledge evidence	Understandability; Resilience	Key judgement, exceptions, holders, redundancy or transfer route.	Interview, observation, successor or peer challenge.	Tacit knowledge may remain incomplete.
Generated Knowledge record	Integrity; Provenance	Tool or process, inputs, transformations, validation, inherited limits.	Reperform or trace a sample.	Vendor opacity or correlated upstream defect.
Authority map	Authority	Create, use, change, challenge and escalation responsibilities and mandate.	Verify role and practical mandate.	Formal chart may differ from practice.

Evidence item	Properties	Minimum content	Verification	Typical limitation
Awareness profile	All	Material dependencies and gaps with State 1, 2 or 3.	Challenge claims.	Self-declaration weakens with consequence.
Access or export test	Resilience	Ability to retrieve or export critical knowledge and rights.	Demonstration.	Technical export may not preserve meaning.
KRO and capability record	Resilience	Capability criterion, target time, unacceptable-consequence point and owner.	Scenario review or recovery test.	A time target without capability is theatre.
Continuity reserve	Resilience	Verified artefacts, access, rights, capability and instructions.	Reconstruction or transfer test.	Archive or escrow alone is insufficient.
Change or migration record	Provenance; Resilience	Old-new mapping, preserved context, validation, rollback or exit.	Sample a migrated case.	Formal completion may hide semantic loss.

18. Pilot gate and open validation questions

This manual is ready for structured validation, not external certification claims.

Validation question	Method concern if observed
Dependency discovery	Assessors repeatedly miss the same material Craft Knowledge, supplier transformation or downstream dependency.
Consequence	The same facts routinely produce widely different profiles without traceable reason.
Evidence sufficiency	Assessors reach opposite sufficiency decisions because the MSE test remains too subjective.
Non-compensability	An essential-property failure is averaged away or hidden.
Outcome model	Assessors disagree on positive versus insufficient or not-assessable reliance for the same evidence.
Revalidation	Late contradictory evidence does not cause reconsideration.
Recovery	Teams pass documentary review but fail operational reconstruction within the KRO and capability criterion.
Gaming resistance	Scope shaving, evidence cloning or ceremonial human review yields an unjustified positive outcome.

Annex A. Minimal assessment record

Record block	Minimum fields
Core summary	Knowledge dependency; consequence; outcome; strengths; uncertainties; actions.
Scope and validity	Object; intended use; boundary; exclusions; date and version; cut-off; revalidation trigger.
Knowledge map	Material Source, Craft and Generated Knowledge; downstream role changes; risk drivers.
Evidence	Evidence IDs; claim linkage; sufficiency rationale; correlation note where material.
Property findings	Each property with state, evidence, limitation and rationale.

Record block	Minimum fields
Challenge	Interviews, observations, samples and tests; independent verification where proportionate.
Modules	Material triggers; module or equivalent-evidence decision; version, status and limitations.
Accountability	Assessor; Accountable Knowledge Authority; operational owner; risk or statutory authority where relevant.
Actions	Owner; action; due condition or date; verification; reassessment requirement.

Annex B. Decision rules at a glance

Decision point	Rule
What is in scope?	Include dependencies that could change consequence, sufficiency, conclusion or action; record exclusions and stopping rationale.
How much evidence?	Collect Minimum Sufficient Evidence for the claim and consequence; do not count documents.
Equivalent evidence?	Only when claim-specific support is comparable and common-source weakness is visible.
Can strong properties offset weak ones?	Not automatically; essential-property material failure caps the outcome.
When use a module?	When a material trigger exists and module or equivalent evidence adds proportionate value or efficiency.
When verify independently?	When consequence, uncertainty, method limits or material dependence make separated challenge proportionate.
What does KRO mean?	Target time paired with an explicit capability acceptance criterion.
When is assurance stale?	When evidence cut-off or revalidation conditions no longer support reliance.
Can a valid assessment be negative?	Yes. Procedural validity and assurance sufficiency are separate.

Stable at the core. Testable at the method. Evidence before confidence.