



ARTISANAL INTELLIGENCE

Government Implementation Guide

v0.6 Release Candidate

English edition

Website publication edition · 17 August 2026

PUBLIC PILOT EDITION · NOT CERTIFICATION

Publication status. This guide supports v0.6 Release Candidate pilots in public administrations. It does not replace legal analysis, public-sector audit, records obligations, procurement rules, data protection, security, AI governance or statutory accountability.

1. What this guide is for

This guide helps a public administration apply Artisanal Intelligence to one bounded knowledge-dependent service, decision process or operational capability. The aim is to reveal and manage material knowledge dependencies without building a parallel bureaucracy.

The first pilot should answer a management question that can lead to action, for example:

- Can this service continue when a key employee retires?
- Can we explain and reconstruct a decision after a system or supplier change?
- Does an AI-supported workflow still preserve authoritative sources and professional judgement?
- Can another competent team recover the required knowledge before citizens or legal deadlines are affected?

2. Select the first pilot

Prefer	Avoid for a first pilot
Recognisable citizen or service outcome.	An entire ministry or heterogeneous portfolio.
Clear process owner and accessible professional experts.	A process with no accountable owner.
At least one visible Knowledge Risk Driver.	A purely theoretical future risk.
Existing operational evidence can be reused.	Massive historical reconstruction before the method is tested.
Management can act on findings.	A showcase selected only because AI is fashionable.
A bounded result is possible within about 90 days.	Organisation-wide certification ambition.

Candidate pilots include permitting, grants or benefits, inspections, crisis response, archival migration, supplier handover, policy interpretation, payroll, case management and AI-assisted administrative drafting.

3. Sponsor and charter

Before evidence collection, agree a short Pilot Charter:

1. **Management problem** - what failure or decision must improve?
2. **Knowledge Assurance Object** - what function, service or decision is in scope?
3. **Intended use** - what must users be able to rely on?
4. **Success criteria** - what observable result will count as useful?
5. **Sponsor and Accountable Knowledge Authority** - who can obtain evidence, challenge conclusions and escalate limitations?

6. **Pilot constraints** - time, people, data protection, confidentiality, procurement and other boundaries.
7. **Explicit non-goals** - no certification, no full legal audit and no new repository unless evidence demonstrates a need.

4. A practical 90-day journey

Stage	Public-sector action	Primary output
Weeks 1-2	Confirm sponsor, object, intended use, consequence screen, stakeholders and exclusions.	Pilot Charter and scope record.
Weeks 2-4	Discover Source, Craft and Generated Knowledge; trace exception paths and dependencies.	Knowledge inventory and Risk Driver Register.
Weeks 4-6	Conduct awareness interviews; index existing evidence; identify correlation and gaps.	Awareness and Evidence Registers.
Weeks 6-8	Assess the five properties; perform proportionate sampling, observation or challenge.	Provisional property findings.
Weeks 8-10	Address material module triggers; test recovery, transfer, provenance or review effectiveness where needed.	Challenge and capability-test records.
Weeks 10-12	Finalise outcome, limitations, actions, validity and pilot evaluation.	One-Page Assessment, Case File and management decision.

The sequence is iterative. A challenge test may reveal a new dependency, change the consequence profile or reopen evidence sufficiency.

5. Governance roles

Role	Core responsibility
Executive sponsor	Authorises the scope, removes barriers and receives material findings.
Accountable Knowledge Authority	Owens the knowledge-assurance position and has practical evidence, challenge and escalation mandate.
Process or service owner	Integrates controls and corrective actions into daily operations.
Professional experts	Surface Craft Knowledge, exceptions, reasoning and context.
Records or archive function	Supports authenticity, metadata, retention, preservation and lawful disposition.
IT, data and AI functions	Supply system, data, model, access, logging, configuration and migration evidence.
Procurement and vendor management	Address exit, access rights, documentation, subcontracting, portability and supplier knowledge.
Assessor	Applies the method, challenges proportionately and reports uncertainty without overclaiming.
Risk, legal, DPO, security or internal audit	Provide specialist or independent challenge where relevant; participation does not create certification.

Do not make the Accountable Knowledge Authority a paper owner. If the role cannot obtain supplier evidence, challenge the assessment or escalate unresolved limitations, governance is incomplete.

6. Reuse existing governance and evidence

- Reuse risk registers, records schedules, continuity plans, case files, change records, procurement files, AI inventories, data-protection artefacts, security reviews and audit evidence where they support the same claim.
- Map controls once to the five assurance properties; do not copy the same evidence into several repositories.
- Use the Artisanal Case File primarily as an index to evidence and reasoning, not as a duplicate archive.
- Treat certificates and standards mappings as scoped evidence, not as automatic Artisanal Intelligence conformance.
- Record common upstream dependencies. Five reports drawn from one stale dataset are one correlated evidence family.

7. Run the core assessment

7.1 Scope and intended use

Define the object, users, outputs, period, boundaries and downstream use. Challenge exclusions. Stop discovery only when additional dependencies would not reasonably change consequence, evidence sufficiency, conclusion or action.

7.2 Consequence profile

Screen citizen impact, rights, legal validity, safety, resources, continuity, scale, reversibility, duration, detectability, automation and concentration. Basic, Enhanced and Critical are provisional pilot profiles, not public ratings.

7.3 Knowledge and awareness map

Identify material Source, Craft and Generated Knowledge. Look specifically for exception handling, local precedent, informal escalation, supplier transformations and outputs that become sources downstream. Classify material dependencies and gaps as known and evidenced, known but incompletely evidenced, or unknown or falsely assumed.

7.4 Minimum Sufficient Evidence

Begin from evidence already used in normal work. Ask for new material only when a concrete assurance question cannot otherwise be answered proportionately. Test relevance, credibility, coverage, freshness, challengeability, limitations and correlation.

7.5 Five-property assessment and challenge

Assess each property separately. Use interviews, observation, samples, scenarios, reconstruction or transfer tests according to consequence and uncertainty. A strong property cannot average away a material weakness in another.

7.6 Outcome, validity and action

Use the Assurance Manual's controlled four-state outcome model. Record evidence cut-off, review point and event triggers. Assign each material action an owner, due condition and verification method.

8. Suppliers, outsourcing and procurement

Public bodies remain accountable for knowledge used in their functions even when a supplier performs the work.

- Identify knowledge that would become unavailable, unintelligible or unchallengeable on supplier exit.
- Secure practical access, export and use rights together with enough context to interpret the material.
- Identify Craft Knowledge held by supplier personnel and subcontractors.
- Require traceability of material transformations and changes.
- Pair a KRO with a capability acceptance criterion for material supplier-dependent functions.
- For Critical dependencies, test exit, transfer or reconstruction rather than relying on clauses or escrow alone.
- In new procurement, specify evidence, handover, continuity and reassessment requirements in proportion to consequence.

9. Records, archiving and migration

Records and archival controls provide powerful evidence for Source Knowledge, Provenance and long-term Understandability. Artisanal Intelligence creates no independent authority to destroy records.

Migration assurance tests more than file survival. It asks whether provenance, metadata, rights, meaning, links, business context and the ability to interpret exceptions survive. A technically successful migration can still represent knowledge loss.

10. AI and automation

- Classify AI as a Knowledge Risk Driver and identify the knowledge it consumes, transforms and produces.
- Preserve other legal and governance obligations; Artisanal Intelligence examines the knowledge-assurance position.
- Record material generated contributions when they affect public decisions, explanations, rules or continuity.
- Protect against circular reuse of generated content and the disappearance of primary sources.
- Treat human review as evidence only when reviewer competence, mandate, time, access and actual challenge are credible.

- When AI is a material trigger, use the AI Handling Module or equivalent evidence and record the decision.

11. Optional Assurance Modules in government

Modules are optional mechanisms, not opt-outs from addressing material uncertainty.

Module	Public-sector use
AI Handling	Test source linkage, review effectiveness, generated contributions and change triggers.
Provenance Verification	Trace multi-system, supplier or archival transformation chains.
Knowledge Recovery	Define capability criterion and KRO; test key-person, supplier or migration recovery.
Automation and Monitoring	Detect when evidence, source, model or dependency change requires reassessment.
Standards Mapping	Reuse existing ISO, NIST, records, continuity or audit evidence at claim level.

When a trigger is material, record module used, equivalent evidence used, or reason not used. Method type - human, automated, algorithmic or AI-assisted - does not determine evidence strength or remove accountability.

12. Management reporting

Management question	Concise dashboard answer
What knowledge-dependent services are exposed?	Objects by consequence profile and main Knowledge Risk Drivers.
Where are material weaknesses?	Property-specific findings, with State 3 highlighted.
Can we continue after loss of person, supplier or system?	Capability criterion, KRO, continuity reserve and latest test result.
What requires a decision?	Insufficient or not-assessable items and high-priority actions with owners.
Is the assurance still current?	Evidence cut-off, next review and active change triggers.
Are we improving?	Trend in open material findings and successful transfer or recovery tests.

Do not reduce management reporting to one composite score. A red status without the affected knowledge, consequence and action is not decision-useful.

13. Public-sector case patterns

Case	Knowledge pattern	Risk pattern	Proportionate application
Permit decision	Law and local rules; exceptions, precedent and proportionality; AI-assisted draft analysis.	Retirement plus summaries replacing source reading.	Enhanced or Critical depending on rights; require source traceability, authority, visible reasoning and successor challenge.

Case	Knowledge pattern	Risk pattern	Proportionate application
Benefits or grants	Eligibility sources; exception handling; triage or recommendation output.	Automation plus direct citizen and legal consequence.	Critical where rights are materially affected; unresolved false assumption normally blocks a positive outcome.
Inspection service	Law and protocols; field judgement and evidential thresholds; risk selection and draft reports.	Outsourced analytics plus experienced inspectors leaving.	Enhanced or Critical; capture exception logic and test transfer to another inspector or team.
Municipal crisis response	Plans and contact data; local relationships and escalation knowledge; situation summaries.	Outage, access loss and key-person dependency.	Critical continuity; KRO may be immediate and the reserve should be tested operationally.

14. Evaluate the pilot

Measure the pilot against the framework's quality test.

Dimension	Pilot question	Useful evidence
Usefulness	Did the method expose a material issue or improve a management decision?	Sponsor decision, accepted action, avoided duplication, user feedback.
Measurability	Were scope, evidence, findings, actions and validity observable?	Completed Case File, linked evidence, test result, action verification.
Reproducibility	Did assessors reach materially compatible positions or explain divergence?	Independent review, calibration log, repeated case or parallel-case result.
Proportionality	Was burden justified by consequence and decision value?	Hours by role, new evidence created, evidence reused, delays and benefits.
Gaming resistance	Did the method resist scope shaving, evidence cloning and ceremonial review?	Challenge findings, late contradictory evidence, negative-case outcome.

15. Common implementation failures

Failure	Corrective response
Starting from an AI inventory instead of a public-service knowledge need.	Reframe the object around the task, service or decision.
Treating document count as assurance.	Return to claim-specific Minimum Sufficient Evidence and operational challenge.
Making a named owner responsible without mandate.	Clarify access, challenge, escalation and risk-acceptance relationships.
Calling optional modules irrelevant.	Record material trigger treatment and accept that a module-free outcome may be negative.
Publishing a positive conclusion without period or cut-off.	Add validity metadata and change triggers; supersede stale claims.
Attempting organisation-wide rollout before pilot learning.	Freeze the core, test one bounded object, measure burden and calibrate.

Annex. Management review agenda

1. Confirm object, intended use, consequence profile and material exclusions.
2. Review the five property conclusions separately.

3. Challenge State 3 items, evidence gaps and common-source dependencies.
4. Review module decisions and results of operational tests.
5. Decide acceptance, corrective action, escalation or temporary restriction of reliance.
6. Confirm owners, due conditions, evidence cut-off and revalidation triggers.
7. Record pilot usefulness, measurability, reproducibility and burden.

Start with one service. Reuse real evidence. Make uncertainty visible. Test whether the knowledge can survive change.