

ArtiProof

Keeping Organisational Knowledge Dependable in the Age of AI

What must an organisation continue to know
in order to act reliably?

Public Whitepaper 1.2 EN

16 September 2026

ArtiProof publication series | KopKen VZW

Informative publication. Research-ready; not yet broadly independently validated. The first independent replication remains pending.

AI-assisted drafting and editorial preparation. No independent human peer review is claimed. Publication and method version numbers are separate.

Executive summary

Artificial intelligence is changing how organisations create, interpret and use knowledge. It can make work faster, more consistent and more scalable. It can also allow an organisation to continue producing answers while gradually losing the ability to explain where those answers come from, who understands them, which judgement still matters and how the work could be recovered if a person, system or supplier disappeared.

ArtiProof is designed to help organisations examine that possibility without assuming that automation necessarily weakens knowledge.

ArtiProof provides a structured way to examine which knowledge important work depends on and what evidence supports relying on it. It can guide an assessment or a peer review of an existing assessment. Critical examination may confirm a bounded conclusion, qualify it, challenge it, or find that the evidence is insufficient. Research must establish how consistently and usefully the method supports those judgements.

The method distinguishes three forms of organisational knowledge: Source Knowledge, Craft Knowledge and Generated Knowledge. It examines them through five assurance properties: Integrity, Provenance, Authority, Understandability and Resilience.

ArtiProof is deliberately not a general management system, a maturity ladder, an AI-safety standard or a documentation programme. It is a focused knowledge lens that can work alongside existing quality, process, risk, governance, HR, continuity, records and technology practices.

Its central design principle is simple: make the basis for reliance explicit, and match the additional review effort to what is at stake. The intended benefit is better-supported judgement, not a guarantee of correctness or an obligation to discover another defect.

A service that works — until the usual answer is challenged

A team completes its daily work successfully. The procedure is available, the system responds and an assistant can explain what to do next. Then an unusual case arrives. Two instructions conflict. The experienced colleague who usually resolves the issue is away.

Nothing about this hypothetical situation proves that the organisation has lost its knowledge. It does, however, expose a question that routine performance alone cannot answer: can the team establish which instruction deserves reliance, explain the uncertainty and choose a defensible next action?

ArtiProof begins with that question. It does not ask organisations to know everything or to distrust every automated answer. It asks them to make the knowledge behind consequential work examinable.

1. Knowledge is operational infrastructure

Every organisation relies on infrastructure that can be seen: buildings, systems, applications, contracts, procedures, organisational structures and data.

It also relies on infrastructure that is much harder to see: knowledge.

A process works because someone understands what a rule means. A public service remains lawful because legislation has been translated into interpretation and practice. A technician recognises when the documented procedure is not enough. A team knows which exception matters. A manager understands which figure can be trusted and which is only an approximation.

This knowledge is rarely stored in one place. Some of it sits in documents and records. Some lives in the judgement of experienced people. Some is embedded in software, analytics and automated rules. Increasingly, some is generated or transformed by artificial intelligence.

That makes organisational knowledge both distributed and fragile.

ArtiProof starts from a basic management question: what does this organisation need to continue to know in order to understand, justify, perform and recover its important work?

2. The efficiency paradox

Digitalisation makes information easier to store, retrieve and transform. AI makes it easier to produce summaries, recommendations, classifications, draft decisions and explanations.

That is valuable. But access to more information does not automatically create stronger organisational knowledge.

An organisation can become faster while becoming less capable of explaining what its systems do. It can automate a task successfully while allowing the human ability to challenge an unusual case to decay. It can retain every document while losing the people who understand why two contradictory documents exist. It can generate plausible answers while becoming less able to distinguish authoritative knowledge from copied or transformed material.

This is the efficiency paradox at the heart of ArtiProof: operational performance may improve while knowledge resilience weakens.

Nothing has to fail immediately. The weakness may remain invisible until a key expert leaves, a supplier changes, a model is challenged, regulation changes, a rare case appears, a system becomes unavailable or an organisation must reconstruct a decision years later.

3. Information is not the same as knowledge

A repository with thousands of documents contains information. Knowing which three documents govern a consequential decision is knowledge.

A database contains data. Understanding what a field means, why an exception exists and when the data should not be trusted is knowledge.

An AI system can produce an answer. Knowing why that answer deserves reliance, who can challenge it and what happens when the system is wrong is an organisational capability.

ArtiProof therefore does not begin by asking whether information exists. It asks whether the organisation can connect consequential work to the knowledge, people, sources, systems, evidence and authority on which it actually depends.

4. Why AI makes the issue urgent

The knowledge problem predates generative AI. Retirement, outsourcing, mergers, restructuring, long-lived legacy systems, staff turnover and supplier dependency have always created knowledge risk.

AI increases the speed and scale of the problem because it can move directly from information to action. Work that once required a person to read, interpret and decide can increasingly be performed through a system that summarises, recommends or generates.

The transition can be gradual: human understands; system assists; system recommends; system becomes the default; human capability is used less often; the organisation becomes dependent before it notices.

ArtiProof does not argue that human work should be preserved for its own sake. It asks a more practical question: which human and organisational capabilities must remain available so that automated work can still be understood, challenged and recovered when consequence requires it?

5. Three types of organisational knowledge

ArtiProof uses three knowledge types to distinguish the roles that sources, practical understanding and system-produced material play in a particular activity. These are functional categories, not a claim that all knowledge fits into mutually exclusive boxes.

Source Knowledge is knowledge grounded in an identifiable source: legislation, policy, contracts, standards, specifications, approved records, manuals, process documentation or validated datasets.

Its main strength is traceability to an explicit reference point. Its weakness is that a source can be current yet misunderstood, authoritative yet ambiguous, approved yet obsolete in practice.

Craft Knowledge is knowledge developed through practice, experience, professional judgement and repeated exposure to real situations. It includes recognising unusual cases, understanding why a procedure exists, knowing when a rule requires interpretation, seeing weak signals and understanding dependencies that are not obvious from formal documentation. Craft Knowledge can be reliable precisely because a professional has learned to judge context. It can also become a dangerous single-person dependency.

Generated Knowledge is knowledge produced, inferred, transformed or synthesised by technological systems. It includes outputs from AI, analytics, rules engines, automated calculations, classifications, predictions and generated summaries. Generated Knowledge is not treated as inherently weaker or stronger than human knowledge. It must be supported by evidence appropriate to its use and consequences.

The labels describe the role of knowledge in the assessed context, not its truth or quality. A generated report may become a source for a later decision; recording that change does not make the report authoritative. Source authenticity and traceability likewise do not establish substantive correctness.

The three forms often interact. A caseworker combines legislation with experience. A model transforms internal data. A system recommendation is reviewed through professional judgement. ArtiProof therefore treats knowledge as a network rather than as isolated categories.

6. Five assurance properties

ArtiProof examines important knowledge through five assurance properties.

Integrity asks whether the knowledge is sufficiently complete and protected against inappropriate alteration, degradation or unnoticed change.

Provenance asks whether the knowledge can be traced to its origin and development. An organisation should know, to a degree proportionate to consequence, where important knowledge came from and how it reached its current form.

Authority asks whether the knowledge is sufficiently authoritative for its intended use. Authority may come from law, an approved source, delegated responsibility, professional competence, an accountable role or another recognised basis.

Understandability asks whether the people who depend on the knowledge can understand it and appropriately challenge it. A correct answer that nobody can interpret or question may still create organisational weakness.

Resilience asks whether the organisation can preserve or recover the knowledge and the capability that depends on it when people, systems, suppliers or circumstances change.

The five properties are deliberately kept separate. A knowledge item may have excellent provenance and poor resilience. It may be authoritative but difficult to understand. It may be complete but concentrated in one person. ArtiProof does not hide those differences inside a single maturity score.

7. Scope is not maturity

One of the most important design choices in ArtiProof is to separate scope from strength.

A knowledge dependency may sit with one person, within an activity, across a process, across an organisation or across a chain of organisations and suppliers. A wider boundary is not automatically better or more mature. The correct scope is the smallest boundary that still captures the material dependency.

If a critical knowledge risk is local, the assessment should remain local. If the process depends on a supplier, external expert, shared platform or public partner, the scope should widen until the dependency is visible.

This avoids treating structural expansion as progress in itself.

ArtiProof therefore does not seek to place organisations on a universal Level 1-to-5 ladder. It asks whether the specific knowledge on which consequential work depends is sufficiently assured for its context.

8. Evidence before paperwork

Governance methods can easily become documentation programmes. ArtiProof is designed to resist that tendency.

The central evidence principle is Minimum Sufficient Evidence: gather enough evidence to support a reasonable assurance conclusion, but do not create material merely because a methodology exists.

Existing organisational evidence should be reused before new artefacts are created. Useful evidence may already exist in skills matrices, process descriptions, HR records, architecture repositories, quality files, audit material, responsibility matrices, training records, continuity plans, decision records, system documentation and operational data.

Reuse is not automatic acceptance.

Every reused artefact should be mapped to the precise conclusion it is intended to support and checked for scope, relevance, currency, ownership and reliability. A skills matrix may be out of date. A process diagram may represent the intended process rather than the real process. A past audit may have tested a different question.

Where the consequence is high or documentary evidence is weak, ArtiProof allows reality testing: explanation, demonstration, walkthrough, interview, peer challenge, exception handling, handover or recovery exercises.

The point is not to prove that paperwork exists. It is to obtain enough evidence to justify a conclusion about knowledge. Minimum Sufficient Evidence is not permission to omit legally required records, retention, reporting or tests.

9. Professional judgement is not a control failure

Knowledge-intensive work creates a difficult design problem. Reliability cannot always be achieved by writing a more detailed procedure.

Professional judgement may be essential because the real world contains ambiguity, exceptions and context that cannot be fully encoded.

ArtiProof does not assume that autonomy is the opposite of assurance. This is a Candidate/experimental implementation pattern: research asks whether making the conditions for professional judgement explicit supports assurance without prescribing every judgement.

Those conditions may include appropriate competence and experience; access to authoritative sources; clear decision rights; explainable reasoning where consequence requires it; peer challenge; transfer to another capable person; and a credible recovery route.

Professional autonomy is never a waiver from assurance. Equally, documentation density is never proof of control.

A useful principle is: manage the dependency and the evidence around craft knowledge; do not automatically proceduralise the craft itself.

10. Consequence determines assurance

Not every knowledge item deserves the same controls.

The consequences of weak knowledge may range from inconvenience to legal error, financial exposure, operational disruption, public-service failure, safety impact or inability to recover an essential capability.

ArtiProof uses consequence-based proportionality. The greater the consequence of knowledge failure, the stronger the evidence and assurance that should be expected. The method asks assessors to match evidence and challenge to the consequences, uncertainty and dependencies of the case. This is a design rule. It does not establish that the resulting assessment will be proportionate in practice; additional effort and decision value must be evaluated.

The target is not maximum maturity. It is sufficient capability.

That distinction matters. Without it, organisations are encouraged to accumulate controls, documents and review layers simply because “higher” appears better. ArtiProof instead asks what level of assurance is justified by the actual consequence, context, rate of change and difficulty of recovery.

11. Assurance must remain alive

A static assessment decays.

Knowledge changes when people leave, systems are replaced, suppliers change, regulation evolves, AI is introduced, an incident exposes a hidden assumption or a once-routine skill stops being practised.

ArtiProof therefore treats assurance conclusions as time-bounded rather than permanently valid.

Each conclusion should state its period, a review point and material changes that could invalidate it, such as departure of a key knowledge holder, a major system or AI change, supplier transition, a material source or regulatory change, a serious incident, a failed handover or an inability to explain or recover a critical capability. A formal event-triggered reassessment mechanism remains Candidate/experimental under controlled method hardening and should not yet be treated as a settled feature of the method. Any reduction in burden or missed changes must be tested.

A proportionate periodic review may be used where appropriate. The research question is whether more formal change-triggered approaches can later reduce burden without increasing missed-change risk or assessor inconsistency.

12. Calibration before complexity

Assessment methods often respond to disagreement by adding more criteria, more scores and more rules.

ArtiProof takes the opposite starting point.

When assessors reach materially different conclusions, the first question is whether the disagreement can be reduced through clearer examples, better evidence guidance, sharper scope and transparent reasoning.

Material disagreement should be recorded before consensus. It may reveal ambiguous criteria, unequal access to evidence, hidden local practices or genuine inconsistency in the organisation.

A reviewer should be able to follow the evidence and reasoning behind an assessment.

Reproducibility is a separate question: do independent assessors reach materially compatible conclusions from the same evidence, scope, consequence framing and guidance? Record their initial judgements before calibration or discussion. Agreement reached afterwards is not evidence of independent agreement.

This principle — calibration before complexity — is especially important for ArtiProof because the method deliberately avoids false precision and composite maturity scores.

12.1 A review can confirm, not merely criticise

Imagine an organisation that has already assessed a consequential service carefully. A reviewer examines the same scope, checks the evidence, questions the assumptions and finds the reliance decision well supported. That can be a useful result. The reviewer need not invent a defect to justify the review.

ArtiProof recognises four possible outcomes: confirm, qualify, challenge and insufficient evidence. Confirmation concerns fitness for a stated use, under stated conditions, during a stated period. Qualification restricts the scope or conditions of reliance. Challenge identifies a materially unsupported conclusion. Insufficient evidence withholds judgement when the available material cannot settle the question.

These are outcomes of examination, not four certification levels. Confirmation is not an instruction to agree. The absence of a new defect, reviewer agreement or a convincing explanation is not, by itself, positive evidence of fitness.

Peer review and confirmation are established practices. ArtiProof does not claim to have invented them. Its candidate contribution is the particular, evidence-reusing structure of the knowledge-fitness review. Whether that structure improves on ordinary competent peer review remains to be established.

12.2 What the reviewer actually examined

A review should disclose the relevant competence and affiliation of the reviewer, prior involvement, conflicts, separation from the original assessment, evidence access, the challenge performed and its limitations. A separate report or a second AI response does not make the underlying evidence independent.

Missing evidence also needs careful handling. A record that was never created differs from one that exists but is unavailable, withheld or no longer retained. Where the reason is unknown, it should remain unknown. A documented absence can explain a limitation; it does not supply the positive support that a confirmation requires.

A reconstructed evidence path helps a reader examine the basis of a conclusion. It does not prove that the conclusion is correct. For AI-assisted work, an account of observable sources, records and transformations is preferable to a plausible but unverified story about a model's hidden reasoning.

13. A bounded assessment flow

The design aims to limit additional work; actual effort depends on the case and must be recorded.

First, define the bounded activity, process, decision, service or capability that matters.

Second, identify the knowledge required to understand, justify, perform and recover it.

Third, determine the relevant mix of Source Knowledge, Craft Knowledge and Generated Knowledge.

Fourth, map the material dependencies: people, sources, systems, teams, suppliers and external parties.

Fifth, reuse existing evidence and identify what remains unproven.

Sixth, assess the relevant knowledge against Integrity, Provenance, Authority, Understandability and Resilience, with the evidence level matched to consequence.

Seventh, identify the smallest sufficient action: reduce a dangerous concentration, clarify authority, preserve a source, capture a critical interpretation, test a handover, improve provenance, add an AI checkpoint, create a recovery route or consciously accept the risk.

The assessment should end in a reasoned decision, not in an inventory. That decision may support the existing reliance judgement, restrict it, reject it, or withhold a conclusion. The five properties remain separately visible: a serious weakness in one is not cancelled by strength in another.

14. The Knowledge Matrix

Many organisations already use skills matrices to understand who can perform which tasks.

The Knowledge Matrix is a Candidate/experimental interface intended to display knowledge dependencies. Its added usefulness and proportionality remain to be tested.

The matrix can connect people to knowledge, processes to knowledge and organisational entities to knowledge. Where dependency crosses organisational boundaries, a chain or ecosystem view can be added.

Its purpose is not to reduce knowledge to a spreadsheet. It is to make concentration, coverage, authority, evidence and dependency visible.

A process may appear well staffed because five people can execute it. The Knowledge Matrix may reveal that only one person understands an exception, one supplier holds the recovery knowledge or no one can explain a legacy rule embedded in software.

That is often the vulnerability that conventional staffing or process data does not show.

15. A one-page management view

Comprehensive thinking does not require comprehensive paperwork.

The evidence base may be detailed, while a management synthesis can remain small. A Candidate/experimental One-Page Knowledge Assurance View is intended to summarise the bounded context, consequential knowledge, principal dependencies, material assurance gaps, recovery position, highest-priority action and direct references to supporting evidence.

The one-page view is not the evidence base and should never become a red-amber-green score for the organisation. It is intended as a traceable synthesis for management discussion and action.

The summary format remains a candidate implementation interface. Its role is to communicate the conclusion, the five-property findings and the evidence boundary; its brevity is not evidence that the underlying assessment was sufficient.

16. AI implementation checkpoints

ArtiProof can be used before and after an AI implementation.

Before implementation, the organisation can ask what knowledge currently makes the process possible, who understands exceptions, where the authoritative sources sit, which human capabilities are consequential and which dependencies already exist.

After implementation, it can ask what knowledge is now generated by the system, which human skills are used less frequently, whether users can still challenge the reasoning, whether provenance is adequate, whether new dependencies have appeared and whether the organisation could continue or recover if the system or supplier became unavailable.

The objective is not to slow automation or require a manual copy of every automated task. It is to make the knowledge consequences visible. An alternative qualified team, another supplier, a different tool or a safe restricted service may provide an appropriate recovery route.

17. Knowledge Pollution

AI dramatically lowers the cost of producing plausible information.

That creates a new organisational risk: generated summaries become sources for new summaries; provisional interpretations are copied into procedures; old and new versions coexist; system-generated material becomes indistinguishable from approved knowledge; and provenance is lost as content is repeatedly transformed.

ArtiProof uses the term Knowledge Pollution for an information environment in which it becomes progressively harder to distinguish what should be trusted, why and on whose authority.

The danger is larger than an individual wrong statement. The organisation can lose the ability to determine which knowledge deserves reliance.

Integrity, Provenance and Authority are therefore increasingly important as the cost of content generation falls.

18. Accountable Knowledge Authority

Important knowledge should not become everyone's responsibility and therefore nobody's responsibility.

ArtiProof uses the concept of an Accountable Knowledge Authority: the person, role or body that can determine whether consequential knowledge is sufficiently authoritative for the organisation to rely on it.

This does not imply a new job title. The authority may already be a process owner, professional lead, legal expert, policy owner, product owner, governance body or other competent role.

The question is simple: who can stand behind this knowledge for this use?

19. What ArtiProof does not try to become

ArtiProof is intentionally bounded.

It is not an AI model, an AI-safety standard, a cybersecurity framework, a records-management system, a business-continuity standard, an HR competency model, a quality-management replacement or a universal management framework.

When an ArtiProof assessment reveals that a knowledge dependency is caused by weak cybersecurity, continuity, records management or another discipline, it should signal the dependency and refer it to the discipline that owns it.

ArtiProof should not absorb every problem it can see.

It also does not claim that a process, decision or body of knowledge is correct merely because the method has been applied. Knowledge assurance supports justified reliance; it does not create truth by declaration.

20. Working alongside established frameworks

ArtiProof's related-framework research examines what can be borrowed from existing practice and where the method could be useful within it. These are different questions. A useful learning mechanism need not justify a new public framework profile, and a conceptual relationship is not demonstrated interoperability.

The publication series develops two selected relationships: a knowledge-management companion for ISO 30401, and a financial-sector application for the EU Digital Operational Resilience Act (DORA). Neither changes the three knowledge types or five properties. The companion papers explain possible interfaces, evidence reuse and boundaries; they do not claim transferred conformity or endorsement.

Other model comparisons remain research inputs. The objective is to learn from control, improvement, delivery and professional practice without importing a second general management system. Existing systems may already answer the relevant questions. In that case, reuse their evidence and avoid duplicate administration.

21. Validation and development status

ArtiProof is an emerging methodology under active development. The validation programme examines independent usability, assessor reproducibility, proportionality, added value beyond existing governance, transferability and organisational outcomes. The first independent replication remains pending. Publication readiness and internal review do not answer these empirical questions.

Its credibility should depend on evidence, reproducible assessment and transparent limitations rather than on the appearance of completeness.

Current research distinguishes several questions that must not be collapsed.

Method validity asks whether the concepts and assessment rules are coherent enough to support the intended conclusion.

Implementation usability asks whether organisations and assessors can apply the method without disproportionate burden. Record gross additional effort, clarification work, evidence reused and the effects on the actual reliance decision. Separate observed benefits from estimates and preserve null, negative and contradictory findings.

Incremental review value asks whether ArtiProof-guided peer review adds anything to ordinary competent peer review, not merely to the absence of review. Match relevant competence, evidence access and conditions; control or report differences in review time. Where feasible, compare competent self-assessment, ordinary peer review and ArtiProof-guided review on sound cases, concealed defects and insufficient-evidence cases.

The measures should include warranted confirmation, false confirmation, appropriate qualification or challenge, correct withholding of judgement, evidence-to-conclusion traceability, initial disagreement and burden. Increased confidence is useful only when it is warranted. A review may be valuable without a novel finding; conversely, an additional finding is not automatically correct or useful.

If ordinary peer review reaches equally well-supported conclusions at equal or lower burden, there is no demonstrated ArtiProof-specific advantage in that setting. That result remains compatible with ArtiProof being a usable structure. Neither case-level agreement nor a successful demonstration validates the method generally.

Sustained integration asks whether the practices remain part of real work rather than becoming a short-lived assessment exercise.

Observable knowledge-assurance effect asks whether using the method actually improves the organisation's ability to identify, understand, trust, preserve or recover consequential knowledge.

A positive result at one stage does not prove the next.

ArtiProof is also testing whether assessors can reach materially similar conclusions from the same evidence, and whether natural-language English and Dutch interfaces preserve the same evidentiary threshold. This whitepaper is informative. It does not constitute certification, accreditation, regulatory approval, academic validation or third-party assurance.

22. Design principles

The method is guided by a small set of disciplines.

Knowledge first: begin with the knowledge required for reliable work.

Reuse before creation: use existing organisational evidence when it genuinely supports the claim.

Minimum Sufficient Evidence: require enough evidence to justify the conclusion, not paperwork for its own sake.

Consequence-based proportionality: stronger consequences require stronger assurance.

Small before complex: begin with a bounded problem and widen only when dependency requires it.

Scope is not maturity: a wider system boundary is not automatically better.

Human capability matters: automation does not remove the need to understand which human judgement remains consequential.

Signal, do not absorb: expose dependencies on other disciplines without turning ArtiProof into a universal framework.

Calibration before complexity: improve examples and evidential guidance before adding scores, criteria or properties.

Independent method: ArtiProof should remain understandable and usable without dependence on a particular vendor, technology or implementation provider.

23. The question ArtiProof ultimately asks

Many organisations know how many employees they have, which systems they run and how much data they store.

The more demanding question is:

What knowledge must we continue to possess in order to act reliably, and how do we know that this knowledge will remain trustworthy, understandable and recoverable when people, technology and circumstances change?

That is the question ArtiProof is designed to make explicit.

Organisations should be able to use artificial intelligence without surrendering the capacity to understand the work for which they remain responsible.

ArtiProof in one sentence

ArtiProof is a knowledge-assurance methodology that structures the assessment and peer review of whether organisational knowledge is fit to be relied upon in a defined context; its consistency and incremental practical value remain subjects of research.

Publication status and references

This informative publication updates English Whitepaper 1.1 with the accepted peer-review positioning and its open implementation questions. Publication 1.2 does not change the normative core or promote the programme beyond its publicly reported v0.8 Empirical Challenge stage.

Research-ready, not yet broadly independently validated. The first independent replication remains pending. The Knowledge Matrix, one-page interface, professional-judgement boundary and formal change-triggered mechanisms discussed here remain candidate implementation patterns where identified. No pilot or independent human review is claimed by this publication.

ArtiProof is a methodology and research programme, not a regulator or certifier. Public availability does not itself grant an open-content licence.

[1] ArtiProof. Methodology and current definitions. <https://artiproof.org/method/>

[2] ArtiProof. Development and validation status. <https://artiproof.org/status/>

[3] ArtiProof. Publication series and English Whitepaper 1.1 baseline. <https://artiproof.org/publications/>

[4] ArtiProof. *ArtiProof and ISO 30401: From Knowledge Management to Knowledge-Fitness Review*. Companion whitepaper 1.1 in this release.

[5] ArtiProof. *ArtiProof and DORA: The Knowledge Behind Digital Operational Resilience*. Companion whitepaper 1.1 in this release.

Method descriptions are programme definitions, not independent effectiveness evidence. Source manuscripts and decision refinements are identified in the internal release record. External sources checked: September 2026.