The Living Knowledge Engine

How governed knowledge becomes trusted learning in every format

A Seven Minute Modules / Group Moovs whitepaper · Version 1.1 (draft)

For CIOs, CTOs, L&D platform owners, learning architects, innovation leaders, knowledge managers, and compliance and risk owners.

In one line

In a world of infinite content, trust is the product.

Executive summary

Most learning content is built as a finished artefact. A module is made, a video is shot, a policy deck is written, and each one is correct on the day it ships. Then the world moves, and each artefact ages on its own, with no shared source and no owner. This is how an organisation ends up with a content graveyard: large stores of material that are unowned, untraceable, ungoverned, and slowly going wrong.

AI does not fix this. AI fills the graveyard faster. Anyone can now generate a plausible module in minutes. The scarce thing is no longer the module. It is the assurance that the module is correct, current, sourced, and owned.

The Living Knowledge Engine is built for that world. It inverts the usual model. Instead of treating each artefact as the product, it treats the knowledge as the product, and the artefacts as views of it. One governed knowledge base holds the sourced, versioned truth. A source registry records where every claim comes from. An Assurance Grade makes that trust visible. From this single base, the same knowledge becomes a module, a video, a podcast, a companion guide, an assessment, an assignment, a manager aid, a job aid, an onboarding sequence, or a grounded AI assistant.

This is not a content factory. A factory produces volume. The engine produces governed, sourced, accountable outputs, each one traceable to the knowledge underneath it and signed off by a named person. AI accelerates the work. It does not own the trust.

The engine is the system beneath Seven Minute Modules. It is not a new brand and not a category above the framework. Seven Minute Modules is what people see and what they remember. The engine is how it stays true. Start with the objective, not the format. The knowledge stays the same. Only the delivery changes.

The product is not the format. The product is the governed knowledge asset.

The old model: content as isolated artefacts

Walk into any established organisation and the pattern is the same. SCORM modules sit in an LMS, last reviewed several reorganisations ago. Videos live on an intranet. PDFs and slide decks fill shared drives. Policy documents sit in a content system. Podcasts sit in a folder. A wiki that began as a knowledge base has become a place where knowledge goes to be forgotten.

Each of these was built once, by someone, from some source, at some time. Each was accurate on its build date. And each now decays on its own clock, because nothing connects them.

That last point is the real problem. The video, the SCORM module, and the policy PDF might all describe the same rule. They were built at different times, by different people, from different sources, so they have already drifted apart. When the rule changes, no one can say which of them is now wrong. There is no shared source to correct, only a scatter of separate artefacts to chase.

We call this a content graveyard: content that is unowned, untraceable, ungoverned, and decaying. The defining trait is not that the content is bad. Much of it was good when it was made. The trait is that nobody owns it, no one can trace its claims, and nothing keeps it current.

AI does not solve this. The same tools that generate one module generate a thousand, each fluent and frozen at the moment it was made. A graveyard you can fill at machine speed is still a graveyard.

The new model: knowledge as the source asset

The fix is to change the unit. Stop treating each artefact as the thing you own. Treat the knowledge as the thing you own, and treat every module, video, and guide as a rendering of it.

Seven parts make that work.

- **A governed knowledge base.** One sourced, versioned foundation that every output draws from, rather than a thousand documents each written from scratch. The knowledge is the asset. The outputs are renderings of it.
- **A source registry.** Every claim in the base is mapped to the source it rests on. This is what makes traceability real, and it is what turns maintenance from a hope into a process. When a source changes, you can see exactly which knowledge depends on it.
- **Versioning.** The knowledge has versions. You can see what changed, when, and why. Nothing changes silently.
- **Evidence mapping.** Each claim is tied to a source and a confidence level: strong where it rests on primary or peer-reviewed work, weaker where it rests on industry or analyst material, with the difference recorded rather than hidden.
- **Ownership.** A named role is accountable for each area of the base. Accountability has an address.
- **A review cadence.** The base is re-reviewed on a schedule, so currency is a routine, not a rescue.
- **Regeneration.** The base is refreshed on a schedule and on change, so it tracks reality instead of drifting away from it.

Together these turn a pile of content into an asset that holds its value. A static library decays. A governed knowledge base compounds.

From source to experience

Here is the full path from a raw source to a published learning experience. It is a loop, not a line. The last step feeds the first.

1. **Source intake.** Bring in the raw material: regulations, standards, peer-reviewed research, official guidance, internal policy, and subject-matter expertise. This is the input, not yet the knowledge.
2. **Source validation.** Check each source. Is it real, authoritative, current, and correctly attributed? Assign a confidence tier. A blog that quotes a blog does not enter the base.

3. **Knowledge extraction.** Pull the claims, rules, definitions, and procedures out of each validated source. Separate what the source says from how it happened to say it.
4. **Knowledge modelling.** Structure the extracted knowledge. Record how claims relate, what depends on what, and which source backs each claim. This is the governed foundation taking shape.
5. **Objective definition.** Decide what a learner should be able to do afterwards. The objective comes before the format. Start with the objective, not the format.
6. **Assurance Grade assignment.** Within Seven Minute Modules, the Assurance Grade is the visible trust label attached to a learning unit. It shows whether the unit is sourced, reviewed, current, traceable, and owned. It is not a decorative badge, and it is not a quality score for style or production value. It is a view of the governance behind the unit: a trust signal for evidence, review, currency, traceability, and ownership. It does not certify that the content is correct forever. It records what the content rests on, when it was last checked, and who stands behind it. At this step the unit's sources, currency, and accountable owner are recorded, so the grade reflects real governance rather than a claim.
7. **Output generation.** Draft the chosen format or formats from the modelled knowledge, with AI assistance for speed. The draft inherits the sources mapped to the knowledge it is built from.
8. **Human review.** A person checks the draft for domain accuracy, instructional quality, and compliance, then signs off. This is the accountability checkpoint, not a final polish.
9. **Publication.** Release the output with its Assurance Grade attached and visible, so a buyer or auditor can see the basis rather than take it on faith.
10. **Monitoring.** Watch the sources for change and the usage for signals: what gets completed, what questions learners ask, where understanding falls short.
11. **Regeneration.** When a source changes or a review falls due, the registry flags the affected outputs. The system prepares updated drafts and routes them through human review before anything is republished.

Build is the first lap. The loop is the product.

The output layer

One governed base, many outputs. The objective decides which one. The same knowledge can become any of the following, each suited to a different moment in a person's work.

The formats vary. The discipline does not. The engine can render knowledge in many shapes, but a Seven Minute Module stays constrained: one objective, a seven-minute envelope, and validated

application. The other formats serve and surround that unit. They do not loosen it.

Format	Best used for	What it adds	Use with care when
Interactive LMS module (SCORM / xAPI)	Trackable, completable learning inside the client's LMS	Completion records, interaction, a graded check	The topic needs human practice or discussion to land
Video	Distribution at scale: intranet, social, embedded	Pace, visual explanation, reach	A learner must follow detailed steps at their own speed (add a job aid)
Podcast / audio	Hands-busy learning: commute, walking, low-screen	Reach into moments a screen cannot	The content is precise and reference-heavy (pair with a written aid)
Companion guide	Reference before, during, or after a module	Something to return to at the point of use	Used as a substitute for the learning itself; it supports, it does not teach
Assessment	Checking and evidencing capability	Proof of understanding, and data for spacing	Treated as the only learning; a test is not a lesson
Workplace assignment	Turning the objective into real work or behaviour	Application, the bridge to transfer	No manager or context supports the follow-through (build that first)
Manager conversation aid	Helping a manager run the learning with their team	Activates the work-environment lever transfer depends on	Used as a script that removes the manager's judgement
Job aid	In-the-moment support at the point of work	Speed and accuracy when it counts	Used as a stand-in for understanding the why
Onboarding sequence	A staged path for a new joiner	Structure, pacing, spacing across the early weeks	It becomes a content dump in week one

Format	Best used for	What it adds	Use with care when
AI assistant / coach	On-demand answers from the governed base	The convenience of "just ask," kept grounded	High-stakes decisions without human supervision (see the note below)
Compliance evidence	Showing an auditor the basis and currency of training	Defensibility: sources, dates, sign-off	Stated as a claim the registry cannot substantiate

Traceability is preserved the same way for every format. Each output is generated from the governed foundation, carries its source-registry mapping, and is published with an Assurance Grade. A video is not a looser artefact than a SCORM module. It is the same governed knowledge, rendered differently, with the same provenance attached. The knowledge stays the same. Only the delivery changes.

A note on the AI assistant. A conversational assistant is the hardest output to keep governed, because it answers in real time rather than being reviewed before release. It earns a place only under firm constraints. It draws answers strictly from the customer's governed knowledge base, not from the open model. Its scope is bounded to what the base covers. It must not answer outside that base, even when the underlying model could. That constraint is deliberate. Every answer is logged and attributable to the sources behind it. When it has no grounded answer, it says so and points to a person, rather than inventing one. Grounding in a verified base cuts error sharply. It does not remove it.[1] So the assistant supports decisions, and high-stakes use stays under human supervision. Handled this way, the thing that worries a content vendor, the wish to just ask an AI, becomes a governed product rather than a new risk.

The engine feeds the human layer. It does not replace it. Several of these formats exist to support people, not to stand in for them. An assignment turns a module into action. A manager conversation aid puts a trusted unit in a manager's hands. An onboarding sequence paces a new joiner's first weeks. Workshops, coaching, roleplay, group discussion, demonstrations, and peer feedback remain essential when the goal is behaviour change, social learning, or judgement. A module can prepare, focus, reinforce, and evidence a learning moment. It cannot replace every human moment. A three-hour classroom session is not the enemy. A three-hour content dump is.

What a client receives

It helps to be concrete. Under this model, an engagement is not a stack of files handed over once. It is five things a client gets and keeps.

- **A governed knowledge base for the domain.** The sourced, versioned foundation for the topics in scope, owned and maintained, not a one-off export.
- **A source registry.** The record of every claim and the source it rests on, so traceability is something you can query, not a promise you have to trust.
- **An Assurance Grade on every published unit.** The visible trust label, showing sources, review status, currency, and the accountable owner, so a buyer or auditor can read the basis in seconds.
- **Reviewed outputs in the formats you need.** Modules, video, podcast, companion guides, assessments, assignments, manager aids, job aids, an onboarding sequence, or a grounded assistant, each generated from the base and signed off by a person. They ship into your existing LMS through standard packaging (SCORM, xAPI, LTI), or are hosted and served for you.
- **A maintenance rhythm.** A scheduled and triggered review cadence that keeps the base current after launch, so the work does not quietly decay.

The first four are where most suppliers stop. The fifth is the one that decides whether the rest stays true.

How it fits, and who owns what

A CIO will eventually ask where this runs, who owns the knowledge, and what happens to confidential material. Left unanswered, those questions get answered by assumption, and the assumptions rarely help. So we answer them plainly.

Where it runs. The operating model does not depend on where it is hosted. It can be delivered as a managed Group Moovs platform, within a customer-controlled environment, or in a hybrid model where security, regulatory, or operational requirements demand it. The method stays the same across all three.

Who owns what. Ownership works the way it does in consultancy and software licensing, with no surprises.

- **The customer owns** their source material, their policies, their internal knowledge, and the customer-specific units built from them.
- **Group Moovs owns** the Seven Minute Modules framework, the governance model, the Assurance Grade model, and the engine processes that produce and maintain the work.

You keep your knowledge. We keep our method.

Who does what. The split is what makes this a managed service, not a tool you have to staff. The customer brings what only the customer has: access to source material, subject-matter expertise, participation in review, the final say on what is published in their name, and the governance decisions about what matters. Group Moovs runs the system: intake, knowledge modelling, generation, the source registry, quality review, and the work of keeping the base current after launch. You decide and approve. We build and maintain.

Your material, and how it is handled. Your confidential sources stay confidential. They are used to build and ground your knowledge base and your outputs, and for nothing else. They are not used to train public or shared AI models. Access controls, auditability, and review workflows follow established information-security practice. Group Moovs works with ISO 27001-aligned and GDPR-aligned processes, with AI governance maturing alongside ISO 42001 and the EU AI Act.

Portability and exit. Because the customer owns the knowledge, the customer can take it with them. On exit, your source material, the contents of your governed knowledge base, and your outputs are yours to export. What stays with us is the method and the engine that produced them, not your knowledge.

The same operating model scales. The engine can support a single course, a curriculum, a compliance library, an onboarding academy, or an enterprise knowledge domain. The operating model stays the same.

Why this changes production economics

The economics change because the unit of work changes. You build the knowledge once and render it many times, rather than building each artefact from scratch.

- **Fewer one-off builds.** A topic is modelled once in the base, then expressed as a module, a video, an assessment, or a guide without restarting the research each time.
- **Reusable foundations.** The same governed knowledge serves every format, so producing the second and third format costs less than building each one from scratch.

- **Faster regeneration.** When a rule changes, the affected outputs are updated from the revised base. You correct the source, not a scatter of copies.
- **One message across formats.** The video and the module say the same thing, because they come from the same source. No reconciliation needed.
- **Less drift.** With a single source, artefacts cannot quietly diverge from one another.
- **Easier compliance updates.** A regulatory change is applied once in the base and carried through to everything that depends on it.

Be clear about what this is not. It is not free content. The expensive, valuable work moves from making content to maintaining knowledge. That is the cost worth paying, because it is the part that keeps the rest true. The cost moves from production to stewardship.

Why this changes quality

Quality stops being a matter of who happened to build a given output.

- **One truth source.** Every output traces back to the same governed base, so there is one version of the answer.
- **Consistent claims.** Formats do not contradict each other, because they share a source.
- **Traceable evidence.** Every substantive claim maps to a source a reviewer can check.
- **Human-reviewed transformation.** A person reviews and signs off each output before release.
- **Visible assurance.** The Assurance Grade shows the basis, so trust is something a buyer can see rather than assume.

There is a hard limit worth naming. Grounding an AI draft in a verified base reduces error sharply, but it does not remove it. A confident, well-formed, wrong answer is still possible even when the right source was available. That is exactly why a human stays accountable, and why the review step is a control, not a courtesy. Trust should be visible, not assumed. This is learning you can stand behind.

Why this changes maintenance

Maintenance is where most content quietly dies, so this is where the engine earns its keep.

- **Source changes trigger update needs.** When a source is revised, retracted, or superseded, the registry knows the base depends on it.

- **Impacted outputs can be identified.** A change becomes a specific instruction: these outputs, and only these, need review. Not a vague worry that something somewhere is now wrong.
- **Review becomes systematic.** Re-review runs on a schedule and on change, rather than as a manual audit nobody finds time for.
- **Content does not become a graveyard.** Because the knowledge is owned and the system tracks decay, nothing ages in silence.

A concrete case makes this visible. If an ISO 27001 control is revised, the registry shows exactly what depends on it: the module that teaches it, the quiz that checks it, the manager aid that supports the conversation, and the job aid people use at their desks. The update becomes a defined list, not a search.

The same view extends across the whole library. The same registry data supports a library-health view of the whole estate: how many units are current, how many are approaching review, how many are awaiting validation, and how many are affected by a source change. The customer is buying portfolio trust, not module trust.

Regeneration does not mean automatic republication. The system identifies the affected knowledge and outputs, prepares updated drafts, and routes them through human review. A person still signs off before anything changes in front of a learner. The machine finds the work. A person approves it.

The pressure is real. Regulatory-intelligence services track a continuous, daily stream of changes across the bodies they follow.[2] A static artefact cannot keep pace with that. A governed base can, because it knows what depends on what. Time is the enemy of a graveyard and the ally of a living system.

The architecture, in text

The whole system, in one view. Read it top to bottom, then follow the loop back to the top.

```
SOURCES
  regulations, standards, peer-reviewed research, official guidance,
  internal policy, subject-matter expertise
    |
    v
GOVERNED KNOWLEDGE BASE
  one sourced, versioned foundation that every output draws from
    |
    v
SOURCE REGISTRY
  every claim mapped to the source it rests on
    |
    v
ASSURANCE GRADE
  sourced, reviewed, current, traceable, owned, made visible
    |
    v
SEVEN MINUTE MODULE DESIGN
  start with the objective: one objective, a seven-minute envelope
    |
    v
OUTPUT FORMATS
  module, SCORM, video, podcast, companion guide, assessment,
  assignment, manager aid, job aid, onboarding, AI assistant
    |
    v
USAGE DATA / FEEDBACK
  completion, questions asked, gaps surfaced, field signals
    |
    v
REGENERATION
  on a schedule and on source change

  [ The loop closes here. Regeneration feeds new and corrected
    knowledge back into the sources and the governed knowledge base,
    so the system tracks reality instead of drifting from it. ]
```

Two things make this an engine rather than a pipeline. The source registry, which lets a change at the top find every output at the bottom. And the loop, which sends usage and corrections back into the base, so the system improves as it runs.

The role of AI

AI does real work in this system, and it is worth being precise about which work.

- It accelerates transformation, turning modelled knowledge into a first draft quickly.
- It helps create variants, adapting a canonical unit into a bespoke one for a specific customer from the same source.
- It helps map source to objective, suggesting how a body of knowledge could become a clear learning aim.
- It generates drafts across formats, from a module to a video script to an assessment.

And it is worth being just as precise about what AI does not do. It does not replace accountability. It does not sign off. It does not own trust.

There is evidence behind that line, not just caution. Models produce errors that look exactly like correct answers, because they are built to generate plausible text, not verified fact. People then over-trust confident automated output and stop checking it. This automation bias affects experts as much as novices and cannot be trained away with instruction alone.[3] The documented fix is accountability: the bias falls when a named person is answerable for the outcome.[4] This is also where the law now points. The EU AI Act framework makes AI literacy an explicit organisational obligation, and high-risk AI systems must be designed for effective human oversight, with attention to the risk of over-reliance.[5] The practical interpretation continues to develop.

So the better the model gets, the more necessary the human checkpoint becomes, not less, because rising trust is the exact condition under which errors slip through. AI makes content abundant. Assurance makes it trustworthy. The engine uses AI for the first and refuses to let AI do the second.

The role of humans

If AI does not own the trust, someone has to. The engine names who, and at which step.

- **Domain review.** Is the content correct? A subject-matter reviewer checks it against the source.
- **Instructional review.** Does it teach, and does it meet the method? A learning reviewer checks the objective, the structure, and the seven-minute discipline.
- **Compliance review.** Does it meet the rules and standards it claims? A compliance reviewer checks alignment and accessibility.

- **Final accountability.** A named role signs off and stands behind the result. That sign-off is recorded in the Assurance Grade.
- **Review cadence.** The same people re-review on schedule and when a source changes, so accountability continues past launch.

This is not humans tidying machine prose. It is humans owning the trust the machine cannot supply. A model can write. It cannot be accountable.

The strategic outcome

Put the pieces together over time and the system does something a library cannot. It gets stronger as it runs.

- Every source mapped makes the base more traceable than it was.
- Every correction improves the base once, and improves every output that depends on it.
- Every regeneration keeps the system tracking reality rather than drifting from it.
- Every review lengthens the trust record, the proof that the knowledge has been stood behind over time.

The advantage is durable because it is structural, not technical. The obstacle is not capability. The obstacle is the operating model. Technology can be copied. Operating models are harder to copy. A content library is built for scale, not for per-customer stewardship. Consulting is built for projects, not for indefinite ownership. Authoring tools sell production capacity, not accountability for the output. Internal initiatives can fund a build, but they struggle to fund a decade of stewardship.

The evidence on that last point is consistent, and we treat it as directional rather than precise. Several analyses converge on one pattern: pilots are buildable, and sustained, governed operation is what fails, on ownership, currency, validation, funding, and accountability.[6] Building the system once is not the hard part. Operating it indefinitely is. That is the part the engine exists to run.

A governed base gets more valuable the longer it runs. A static library does the opposite, losing a little value every month it sits. The gap widens on its own.

What this means in practice

Strip away the theory and the method has a simple shape. For any domain we take on, there is:

- **One governed source base** for the topic or domain, not scattered documents.
- **One registry** of claims and the sources they rest on.
- **One Assurance Grade** on every published unit.
- **One human sign-off chain** that ends with a named, accountable owner.
- **One maintenance rhythm** that keeps it current after launch.

One of each, applied the same way every time. That is what turns a good idea into a method you can buy, audit, and rely on.

Conclusion

The old world made content slow and expensive, so the market learned to mass-produce it and reuse it. AI ended that world. Content is abundant now, and abundance moved the value to trust.

The Living Knowledge Engine is built for that shift. It treats knowledge as the asset and formats as renderings of it. It records where every claim comes from, makes that basis visible through an Assurance Grade, and keeps the knowledge current through a source registry and regeneration. It uses AI to move fast and a named human to stay accountable. And it produces the thing people actually use: trusted, traceable learning, seven minutes at a time.

Most organisations already own the knowledge they need. The challenge is making it governable, traceable, current, and reusable.

Seven Minute Modules is the visible form of all this. The engine is how it stays honest. In a world of infinite content, trust is the product.

A static library depreciates. A governed knowledge engine compounds.

Notes

Editor's Notes

Internal. Not part of the published paper.

Where this sits. This is document 3 in the 7MM family, the system and operating-model paper. It pairs with the strategy paper (*Trusted, Traceable Learning*) and the methodology paper (*Seven Minutes by Design*). It deliberately does not re-prove the seven-minute unit; it treats that as settled in the methodology paper and refers to it. Audience is technical and executive leaders, so the register is system-oriented, with one text diagram and two tables doing structural work.

Positioning guard. The Living Knowledge Engine is framed throughout as the system beneath Seven Minute Modules, not a parent category and not a new brand. Seven Minute Modules stays the framework and the thing the reader remembers. This was a real risk for this document, because the engine has an evocative name; the executive summary and conclusion both state the relationship explicitly so the brand is not demoted.

Not a content factory. The AI and human split is kept explicit and is the spine of the differentiation. AI accelerates transformation; assurance and accountability are human-owned. The trust chain (confident error, automation bias, accountability, the EU AI Act) carries that argument without overclaiming.

Evidence handling. Strong claims are cited plainly (automation bias, accountability, the EU AI Act as primary law). Directional claims are flagged as directional in text and in the notes: the Thomson Reuters ~200-a-day figure as an industry benchmark, and the year-two convergence (including the contested MIT NANDA figure) as a pattern rather than a precise number. The grounding claim now rests on two peer-reviewed sources, Shuster et al. (2021) for the reduction and Niu et al. (2024), RAGTruth, for the residual error, with no specific percentage stated since reported figures vary. No banned statistics, no attention-span myth, no guarantee language.

Verify before publication. Re-pull the Gartner, MIT NANDA, RAND, and S&P figures, since these update. Confirm EU AI Act Article 14 and Article 4 references against the consolidated text. Confirm the Thomson Reuters figure. Confirm all citation details against the master bibliography once it exists. Confirm the hosting model, the IP ownership terms, and the data-handling commitment that customer material is not used to train public or shared models, against the signed DPA and vendor agreements.

Cognitive-load pass applied. Em-dashes removed throughout (commas, full stops, colons, parentheses instead). Sentences kept short, most under about 25 words. Abstract-noun stacks broken into subject-verb-object. Each major section closes on a short, plain line (for example, "A graveyard you can fill at machine speed is still a graveyard," "The cost moves from production to stewardship," "Time is the enemy of a graveyard and the ally of a living system," "A model can write. It cannot be accountable.").

Concreteness pass (this revision). Added "What a client receives" (the five deliverables, with delivery via SCORM, xAPI, LTI, or hosting) and "What this means in practice" (the one-of-each shape), so the paper reads as a productised method, not only a model. Added a worked ISO 27001 example to the maintenance section. Clarified that regeneration routes drafts through human review rather than auto-publishing, in both the lifecycle and the maintenance section. Softened the economics language away from "small/cheap" to "costs less than building each one from scratch." Added an explicit scope boundary to the AI assistant ("must not answer outside that base, even when the underlying model could"). Reinforced the Seven Minute Module discipline in the output layer. Expanded the Assurance Grade definition (not a style or production score; does not certify content correct forever). Trimmed the strategic-outcome section to remove a repeated list and reserve the word "compounds" for its two canonical landings. Replaced the broad RAG note with two peer-reviewed sources.

Fit-and-ownership pass (this revision). Added "How it fits, and who owns what" after "What a client receives," to pre-empt the questions a CIO or CISO will ask before procurement does. It covers deployment (deployment-neutral: a managed Group Moovs platform, a customer-controlled environment, or a hybrid, with no hosting preference asserted, so the paper sells the model rather than the hosting), the IP split (customer owns sources, policies, internal knowledge, and customer-specific outputs; Group Moovs owns the framework, governance model, Assurance Grade model, and engine processes), the division of labour (framed as a managed service, not a tool to staff), data handling and a short security posture (confidential sources not used to train public or shared models; ISO 27001 and GDPR-aligned; AI governance maturing toward ISO 42001 and the EU AI Act), and portability and exit. Added a library-health paragraph to the maintenance section to make the point that the customer buys portfolio trust, not module trust. Added a single scale line that expands the frame from one course to an enterprise knowledge domain. The section is deliberately kept at operating-model level, not solution-architecture depth, and the hosting, IP, and data-training claims are flagged above for confirmation against the signed agreements.

Final pass (this revision). Made deployment neutral by removing the SaaS-as-default framing, so the paper sells the operating model, not a hosting preference. Ran a light de-duplication of "governed knowledge base," which appeared 11 times: a few connective uses were changed to "governed foundation" or "knowledge base," while the term was kept as the anchor at the definition, the canonical landing line ("A governed knowledge base compounds"), the diagram label, the AI-assistant constraint, and the procurement deliverable. Two suggested synonyms were deliberately not used: "source base" (it collides with the source registry) and "canonical knowledge asset" (it collides with the canonical-versus-bespoke distinction). Added a value line near the conclusion: "Most organisations already own the knowledge they need. The challenge is making it governable, traceable, current, and reusable."

Propagation pass (6 June 2026, v1.1). EU AI Act wording aligned to the register's framework framing, with the pending simplification package noted; the regulatory-pace and analyst figures moved to the evidence register per the Appendix B rule; the library-health view reframed as a capability of the registry data pending build confirmation; the most debated analyst source un-named in the notes; the security-process sentence reworded to "works with ISO 27001-aligned and GDPR-aligned processes" pending certificate confirmation.

1. On grounding reducing error: Shuster, K., Poff, S., Chen, M., Kiela, D., & Weston, J. (2021). Retrieval Augmentation Reduces Hallucination in Conversation. *Findings of the Association for Computational Linguistics: EMNLP 2021*, 3784–3803. On grounding not removing it: Niu, C., Wu, Y., Zhu, J., Xu, S., Shum, K., Zhong, R., Song, J., & Zhang, T. (2024). RAGTruth: A Hallucination Corpus for Developing Trustworthy Retrieval-Augmented Language Models. *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 10862–10878. RAGTruth finds that, even with retrieval, models can still produce claims that are unsupported by or contradict the retrieved sources. No single error-reduction percentage is stated here, as reported figures vary by study and setup. ↩
2. Thomson Reuters Regulatory Intelligence, *Cost of Compliance* reporting. Industry source; volumes vary by year and coverage. The current figure is held in the 7MM evidence register (009) and re-verified before client use. ↩
3. Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: an attentional integration. *Human Factors*, 52(3), 381–410. ↩
4. Skitka, L. J., Mosier, K., & Burdick, M. D. (2000). Accountability and automation bias. *International Journal of Human-Computer Studies*, 52(4), 701–717. ↩
5. Regulation (EU) 2024/1689 (the EU AI Act), Article 14 (human oversight) and Article 4 (AI literacy). Article 4, the Article 3(56) definition, and the 2 February 2025 application date verified 6 June 2026 against Commission guidance (see 009); Article 14 wording to be confirmed on EUR-Lex at the next pass. A pending AI Act simplification package may affect timing, implementation details, and guidance. The consolidated text is tracked before any obligation is quoted precisely. ↩
6. Convergent and directional, to be read as a pattern rather than precise figures: Gartner (2024) forecast on generative-AI projects abandoned after proof of concept; an MIT-affiliated enterprise study (2025), methodology publicly debated, treat as directional; RAND (2024), *The Root Causes of Failure for Artificial Intelligence Projects*; S&P Global Market Intelligence / 451 Research (2025), *Voice of the Enterprise: AI & Machine Learning*. ↩