

For maintenance, operations and executive leaders
in asset-intensive organisations

Your most experienced technician is leaving.

How much of your maintenance capability goes with them?

Preserving equipment knowledge with AI, from fault histories and
schematics to the reasoning behind maintenance decisions.

Capture experience. Connect the evidence.
Keep expertise usable.

The records may stay. The reasoning can leave.

An asset trips again. The work-order history records earlier repairs. The schematic shows the circuit. An experienced technician knows which past incident is relevant, which modification changed the behaviour and which apparent fault is a distraction.

That is the capability at risk when a key person leaves. The organisation may retain its files while losing the ability to connect them quickly and correctly. This opening scenario is illustrative.

What leaders are trying to preserve

For maintenance, the aim is to make equipment knowledge usable beyond the people who accumulated it. For operations, it is continuity across shifts and sites. For executives, it is less dependence on individuals and a clearer path to developing the next generation of technicians.

A September 2026 study by Deloitte and The Manufacturing Institute links technician shortages and skills gaps with downtime and constrained production capacity. It examines how AI could help people develop and apply expertise in daily work. [1]

Where AI fits

An AI assistant can retrieve relevant records and help explain their connections. Its usefulness depends on whether those records describe the right equipment, preserve the expert's reasoning and remain current. Access to an answer does not establish competence to act on it.

This paper examines an equipment-centred application, related industrial implementations and current research. It sets out what to capture, where assistance should stop and how to test whether the system makes knowledge more usable.

The investment question

Maintenance

Can another qualified technician find and understand the evidence?

Operations

Does knowledge remain usable when the expert is unavailable?

Leadership

What capability will remain with the organisation when people move on?

Organise the knowledge around the equipment.

The implementation described for this paper brought together equipment failure logs, work-order history, electrical and hydraulic schematics, protection and trip settings, and codified technician knowledge. It organised that material around equipment and fault zones. [P]

Give scattered sources a shared context

A fault zone provides a defined area of the equipment around which related information can be organised. The purpose is to connect a question with the relevant equipment history and engineering information, rather than search an undifferentiated collection of documents.

This matters when an explanation spans several sources. The work order may record an intervention, a drawing may show the relevant relationship, and the technician's account may explain why the earlier decision was made. Each contributes something different.

Make experience retrievable

The application used retrieval-augmented generation, or RAG: relevant information is retrieved from a maintained knowledge collection and supplied to the model to help it answer. The technician's recorded knowledge becomes part of that collection alongside the technical records.

Keep boundaries in the design

The described guardrails included qualification requirements, including high-voltage qualification where applicable, and a restriction against generating isolation instructions. Those are design boundaries. Whether they were consistently enforced would require evaluation.

The operational test is whether another qualified person can use those connections to understand the equipment without relying on the original expert being available.

[P] Practitioner account supplied for this whitepaper. Implementation description, not an independently evaluated performance case.

Sources brought together

Equipment context

Equipment identity and fault zones.

Operating history

Failure logs and work orders.

Engineering records

Electrical and hydraulic schematics; protection and trip settings.

Human knowledge

Recorded technician explanations and experience.

Record why the decision made sense.

A repair record can tell the next person what was changed without explaining why. Useful knowledge capture preserves the reasoning that connects the symptom, the evidence and the decision.

Start with a completed case

Choose a recurring fault or an incident for which the team repeatedly consults one person. Review it with the technician and a second qualified reviewer. The aim is to make the reasoning understandable to someone who was not there.

Ask what alternatives were considered, which evidence ruled them out and what remained uncertain. Distinguish an approved method from a local workaround. A plausible account should not become approved guidance simply because it has been recorded.

Capture the connections explicitly

Link the account to the asset, relevant fault zone, work order and applicable drawing revision. Record the operating circumstances and the outcome, including whether the fault returned. A lesson without its limits can be applied to the wrong equipment or conditions.

Keep the knowledge reviewable

Give each approved entry an owner, a review status and links to its evidence. Retain disagreements and unresolved questions for review. Update or retire the entry when the equipment, procedure or underlying evidence changes.

Build capture and review into the completion of maintenance work. A one-off interview creates a snapshot; a maintained record can reflect what the team learns next.

A practical capture record

Situation

Asset, configuration, symptom and operating conditions.

Reasoning

Evidence used, alternatives rejected and remaining uncertainty.

Decision and result

What was approved, what happened and what was learned.

Limits

Where the lesson applies, who reviewed it and when it expires.

The approach is already appearing in industrial products.

Two commercial examples show how suppliers are connecting maintenance information with human experience. Both make earlier knowledge available within the context of a current maintenance question.

Siemens Senseye: reusing maintenance cases

Siemens' February 2024 announcement described generative AI features that find similar past cases and solutions, use maintenance notes and process information from different maintenance systems. It named BlueScope Steel as an early adopter and described knowledge sharing across teams as an objective. [2]

The relevance is specific: an earlier case becomes more useful when the next team can find it and understand how it relates to the current issue. The announcement is a supplier account, not an independent measurement of knowledge retention.

ABB Genix Copilot: combining operating evidence

A November 2025 Microsoft customer story describes ABB combining live equipment data, error and maintenance logs, manuals and expert knowledge stores to support troubleshooting. It also describes an industrial deployment at Birla Opus Paints, although those broader platform results cannot be attributed solely to knowledge capture. [3]

The practical comparison is how the product connects current evidence with the relevant history. A buyer should ask to see those connections on their own equipment, including incomplete and conflicting records.

What to take from these examples

Evaluate how a solution retrieves relevant experience and ties it to current evidence. A conversational interface is useful only if the information behind the answer is appropriate to the equipment and the decision.

Commercial sources describe supplier capabilities and customer accounts. This paper does not adopt their headline savings claims as independently established results.

Ask for the evidence

A real case

Which records informed the answer?

A relevant match

Why does the earlier event apply to this asset?

A clear boundary

What remains uncertain or requires expert review?

Useful reasoning needs the right evidence.

IBM Research: checking conclusions against evidence

IBM's July 2026 publication summary describes Condition Insight Agent, combining work orders, operational indicators and structured failure knowledge. It reports a rule-based verification process intended to suppress unsupported conclusions and describes cases from production maintenance systems. [4]

This supports investigating evidence checks as part of an assistant, alongside retrieval. The public summary does not provide enough detail to establish a general error rate or a maintenance return on investment.

MindRAG: connecting signals, annotations and history

The June 2025 MindRAG preprint structures condition-monitoring information, annotations and maintenance work orders for retrieval and AI-assisted analysis. Preliminary evaluation involved an experienced analyst. The authors identify remaining work in retrieval, adaptability and explanation. [5]

Its relevance is the combination of machine context and recorded experience. Preliminary research results should not be presented as proof of reliable performance across all assets or sites.

Multimodal research: using drawings as well as text

A November 2025 study in Advanced Engineering Informatics combines text and schematic features to construct equipment-maintenance knowledge graphs. It evaluates extraction of entities and relationships using an automotive maintenance dataset. [6]

That is relevant to electrical and hydraulic schematics, where relationships can matter as much as words. The reported evaluation concerns knowledge extraction and graph reasoning, not a demonstrated reduction in industrial downtime.

Together, these studies provide reasons to test structured, evidence-based assistance. They do not remove the need to evaluate it with your own equipment and qualified reviewers.

Three evidence levels

Publication summary

A reported approach and deployment claim, with limited public evaluation detail.

Research evaluation

Performance on specified data and tasks.

Operational result

Measured benefit in a defined working environment.

Better AI can make your equipment knowledge more useful.

A general model may become better at reading diagrams, comparing records and explaining technical relationships. It still needs access to the correct history, configuration and approved knowledge for your equipment.

Preserve context that the model does not already have

The investment is in maintaining that context: asset relationships, reviewed technician accounts, source revisions and records of what happened. A more capable model can potentially make better use of it. The benefit remains conditional on the quality of the sources and the model's performance on the task.

Treat an upgrade as something to test

Use the same expert-reviewed questions and the same source collection to compare model versions. Include ambiguous symptoms, similar assets with different configurations, conflicting records and requests outside the assistant's allowed scope.

Check whether the new version identifies the correct sources, explains their relevance, recognises missing information and respects restrictions. Track any improvement alongside regressions, response time and operating cost.

Do not expect a model upgrade to repair missing knowledge

A stronger model cannot reliably reconstruct an undocumented modification or confirm that an obsolete drawing still applies. If ingestion loses a diagram label or retrieves the wrong asset, better language generation does not resolve the underlying problem.

NIST's 2026 work on deployed AI highlights the need to monitor functionality and human interaction, among other areas. For a maintenance assistant, continued evaluation should cover both model changes and changes in the equipment knowledge it uses. [7]

What may improve

Interpretation

Understanding diagrams and technical language.

Comparison

Connecting a present question with relevant earlier cases.

Explanation

Showing what the evidence supports and what is missing.

What remains essential

Current records, expert review and enforceable permissions.

Knowledge access is not permission to act.

Making an explanation easier to obtain does not qualify a person to perform the work. A useful assistant needs explicit limits on what it can reveal, recommend and initiate.

Separate explanation from execution

The practitioner application described qualification-based guardrails and excluded generated isolation instructions. In a production design, role and qualification checks should be enforced by the application and its access controls, rather than depend only on the model following a prompt.

The assistant should identify restricted requests and direct users to the responsible qualified personnel and approved site procedures. It should not invent an alternative procedure when the required information or authority is missing.

Make the evidence visible

Answers should identify the equipment, the source records and the applicable revisions. Separate a documented fact from an inference. If records conflict, show the conflict for review instead of silently choosing the more convenient explanation.

Assign someone to maintain the knowledge

Engineering changes, revised settings and new maintenance outcomes can make an earlier answer obsolete. Give an accountable owner responsibility for review, withdrawal and updates. Preserve the ability to see which sources supported an earlier response.

Test where the system must stop

Evaluation should include wrong-asset questions, obsolete documents, absent qualification, unsupported conclusions and requests for excluded instructions. Record whether the system refuses, asks for clarification or escalates as intended.

Review boundary failures alongside successful answers. A helpful-sounding response to an excluded request is still a control failure; it needs investigation before wider use.

Assign clear responsibilities

Knowledge owner

Maintains records and review status.

Technical reviewer

Checks equipment-specific accuracy and applicability.

System owner

Enforces access and manages changes.

Operational owner

Defines scope and the route for escalation.

Prove the capability on a bounded maintenance problem.

Start with one equipment family and a recurring information problem. Select cases for which reliable source material and experienced reviewers are available. The first goal is to make useful knowledge easier to find and interpret.

Establish the baseline

Record how people currently find the relevant history, drawings and expert advice. Use representative questions and note the time, missing sources and escalations involved. Measure the full task, including checking the answer.

Compare assistance with current practice

Have qualified reviewers assess the same cases with and without the assistant. Include typical cases and difficult exceptions. Keep the equipment, source versions and scoring criteria consistent so that differences are interpretable.

Judge usefulness and correctness together

A fast answer about the wrong asset is not an improvement. Measure source accuracy, technical acceptability, appropriate escalation and review effort alongside time saved. Do not claim reduced downtime from faster information retrieval alone.

Budget for maintaining the knowledge

Include expert capture time, source preparation, integration, review, permissions, model usage and ongoing updates. Preserve exportable records and links so that accumulated knowledge can survive a change of model or supplier.

Expand only when the evidence supports it

Agree acceptance criteria before the trial. Resolve critical boundary failures before widening access. Extend to another asset family or site only after checking its records, terminology and operating context.

The objective is continuity of maintenance capability: another qualified person can find the evidence, understand the reasoning and recognise when expert help is still required.

A compact scorecard

Evidence

Correct asset, source and revision.

Answer quality

Accepted by a qualified reviewer, with uncertainty made clear.

Effort

Time to a checked answer and demand on experts.

Boundaries

Correct handling of restricted and unsupported requests.

Evidence behind the paper.

The paper combines research, commercial implementation accounts and a practitioner description. These are different kinds of evidence. The practical recommendations are NextGenMaint's synthesis, not measured results from a single deployment.

[1] Deloitte and The Manufacturing Institute

Expanding the skilled manufacturing workforce with AI. 9 September 2026. [Read source](#)

Workforce context and proposed uses of AI; not a trial of the practitioner application.

[2] Siemens

Siemens elevates predictive maintenance with innovative generative AI. 7 February 2024. [Read source](#)

Supplier announcement naming BlueScope as an early adopter and describing case retrieval and knowledge sharing.

[3] Microsoft / ABB

ABB leads shift to autonomous operations with Genix Industrial AI platform on Azure. 18 November 2025. [Read source](#)

Commercial implementation account. Broader platform results are not isolated measures of knowledge preservation.

[4] O'Donncha and colleagues / IBM Research

Evidence-Driven Reasoning for Industrial Maintenance Using Heterogeneous Data. Listed as ACL 2026; 2 July 2026. [Read source](#)

Public publication summary consulted. Production use and verification approach are author-reported.

Scope and interpretation.

[5] Löwenmark and colleagues

Agent-based Condition Monitoring Assistance with Multimodal Industrial Database Retrieval Augmented Generation. arXiv:2506.09247, version 1, 10 June 2025. [Read source](#)

MindRAG research preprint; preliminary analyst-assisted evaluation.

[6] Zhang and colleagues

A knowledge graphs construction method enhanced by multimodal large language model for industrial equipment operation and maintenance. Advanced Engineering Informatics 68, 103705, November 2025. [Read source](#)

Publisher abstract consulted. Automotive data used to evaluate knowledge extraction and graph reasoning; full-text findings not independently assessed.

[7] NIST

New Report: Challenges to the Monitoring of Deployed AI Systems. 9 March 2026, updated 18 March 2026. [Read source](#)

Overview of NIST AI 800-4 and monitoring challenges. Used for the case for continued evaluation, not as certification guidance.

[P] Practitioner implementation

Implementation details were supplied by the developer for this publication. They cover equipment fault zones, failure and work-order history, schematic ingestion, protection and trip settings, technician knowledge and intended guardrails. No customer identity, deployment scale or evaluation results were provided. The case is not independently audited.

How to use this paper

Knowledge-capture practices, controls and evaluation steps are recommendations, not additional claims about the practitioner implementation. Use them to frame a bounded evaluation with your maintenance, engineering and operations teams. The examples do not replace approved technical procedures or authorisation. Claims about improved use of knowledge as models advance are hypotheses to test, not a guarantee that each model upgrade improves the system.

Prepared October 2026. Public sources reviewed 4 October 2026. No vendor endorsement or general performance guarantee is implied.