

Your AI Can Access the Data. Is the Data Ready for AI?

Enterprise AI Adoption Gap | Eldris Inc | October 7, 2026

A decision guide for deploying an internal policy and knowledge assistant

© 2026 Eldris Inc. All rights reserved.

Executive brief

An enterprise can connect an AI assistant to its information without establishing that the information is fit to answer employees' questions. Access is one requirement. Appropriate access, authoritative sources, applicable policies, successful retrieval, and dependable answers are separate requirements.

That distinction changes the investment decision. A disappointing answer may require better retrieval rather than more data. A persuasive but outdated answer may require content remediation rather than a stronger model. An answer that exposes sensitive information may reflect an existing permission problem, an application authorization failure, or both. Treating all three as a generic AI quality problem sends money toward the wrong intervention.

Microsoft documents that Copilot honors existing access controls and separately prescribes remediation of oversharing. AWS makes a related distinction for custom knowledge bases: filtering against supplied access lists does not itself authenticate the person making the request. These are different architectures, but both make the enterprise's information and identity arrangements consequential. [1, 2, 3]

The business decision: launch with a curated collection, remediate specific information and access problems, or expand to a broader collection after demonstrating incremental value and acceptable exposure.

The bounded question: for an internal assistant answering employee questions about company policies and procedures, what evidence should determine which of those three paths to take?

The population: employees in a defined business unit and the deployment, security, and information owners accountable for the assistant. Country, employment category, and role must be represented where they change the applicable answer.

Eldris recommends choosing the smallest maintained collection that passes the workflow requirements, then testing expansion against identifiable unanswered needs. This is a decision principle, not an empirical finding that smaller collections always perform better. Broader access can be the right choice when it adds necessary evidence and its controls hold.

What the evidence establishes

This report synthesizes primary product documentation and published research. It does not report an Eldris survey, deployment audit, or benchmark. The framework, examples, and recommended evaluation design are Eldris analysis. Examples are illustrative and contain no measured outcomes.

Evidence	What it supports	What it does not establish
Microsoft Copilot data-protection architecture [1]	Existing authorization and protection settings shape what the assistant can reference.	That existing permissions match today's intended business access.
Microsoft's deployment guidance [2]	Oversharing, ownership, and obsolete content warrant explicit remediation.	How prevalent these problems are across enterprises, or their financial impact.
AWS custom knowledge-base access controls [3]	Applications must supply verified identity; ACL filtering has defined limits.	That every managed or custom assistant has the same permission behavior.
Microsoft's retrieval and response evaluation guidance [4, 5]	Retrieval relevance, groundedness, completeness, and correctness answer different questions.	A universal score that proves readiness for every workflow.
Liu and colleagues' long-context experiments [7]	The tested models' ability to use information depended on its position in context.	That current enterprise assistants have the same failure rate or behavior.

The evidence supports a diagnostic approach: investigate access, information, retrieval, and response behavior separately. It does not support a claim that information readiness is the largest adoption barrier or that any specific deployment strategy guarantees return on investment.

The documentation is especially detailed for Microsoft 365 and AWS. The management questions travel across platforms; the documented control behavior does not. Verify the actual product, connector, license, configuration, and data path before applying a recommendation.

Access is one boundary, not the whole decision

There are at least three distinct boundaries to examine.

The authorization boundary determines which information this user is allowed to access through this application. It needs an authenticated identity and enforced permissions across retrieval and output paths.

The discovery boundary determines which permitted information the assistant can locate or draw into an answer. Making content harder to discover can reduce exposure, but it is not necessarily equivalent to removing access.


The authority boundary determines which information the organization recognizes as governing the question. A readable document may be a draft, an expired policy, a working note, or guidance for a different country.

Microsoft describes Restricted Content Discovery as a temporary measure while organizations review access. It does not change permissions or prevent existing direct access, and its scope has exceptions. Microsoft also warns that excessive restriction can weaken answer completeness and relevance. The decision therefore involves both exposure and usefulness; disabling discovery is not a complete permission repair. [6]

For custom sources, AWS documents that access lists are supplied by the customer and are not verified in real time against a native permission system. Its filtering relies on identity context supplied by the application. The enterprise should verify how permission changes reach the index and how the application prevents users from substituting someone else's identity. [3]

Access is not the same as readiness

Three independent boundaries determine whether an answer belongs in the service.



Passing one check does not establish that the other two pass.

ELDRIS / ENTERPRISE AI ADOPTION GAP

© Eldris Inc

Figure 1. Three boundaries, three checks. Eldris conceptual framework. A system can pass one check and fail another; the diagram represents no measured adoption rates.

The practical implication is to avoid accepting either “it respects permissions” or “we excluded those sites” as a complete readiness case. Ask which boundary the control addresses, which paths it covers, and what evidence shows that it works for the intended population.

Make the collection answerable

A collection is answerable when the organization can establish the governing answer for the questions the assistant is expected to handle. Eldris proposes five checks. These are operational questions, not a validated maturity scale.

Check	Evidence to request	Consequence if unresolved
Coverage	Required questions mapped to approved sources, including exceptions.	Narrow the promise or obtain missing information.
Authority	A named owner identifies the governing version and how conflicts are resolved.	Correct ownership and conflicting guidance before presenting a definitive answer.
Applicability	Country, role, employment category, effective date, and other relevant conditions.	Collect necessary context or route the question to a person.
Appropriate access	Intended audiences compared with effective access and application behavior.	Fix excessive access and validate restrictions.
Maintenance	An owner, update process, and tested route for changes and revocations.	Establish a maintained service before expanding reliance.

Consider an employee asking, “Can I expense this hotel?” A travel policy may be accurate for one country and wrong for another. A reimbursement threshold may have changed, while an older presentation remains easier to retrieve. A manager’s informal note may describe an exception that does not apply to the employee.

The assistant needs to identify the applicable policy and the employee context required to interpret it. If those facts are missing, requesting clarification or providing an appropriate escalation can be the successful outcome. A confident answer is not automatically the useful outcome.

For multinational enterprises, broad availability of documents should not be mistaken for a single global policy. Country and business-unit exceptions need explicit treatment in the source collection and evaluation. Translating a policy does not establish that it applies in another jurisdiction.

A curated collection also has costs. It can omit necessary exceptions, duplicate records, and create a second maintenance burden. Where possible, curate access to maintained authoritative sources rather than create an unmanaged copy. The preferred collection is one the enterprise can maintain and evaluate, not merely one that is small.

Diagnose the failure before funding the fix

In an assistant that retrieves company documents before producing an answer, the failure can occur at several stages. Eldris recommends recording the question, user context, eligible sources, retrieved passages, answer, and reviewer judgment together. Protect these evaluation records according to their sensitivity.

Observed failure	Competing explanations	First investigation
A necessary rule is missing.	Source absent; permission blocked it; retrieval missed it; answer omitted it.	Locate the approved rule, then inspect each stage.
An outdated rule is given.	Obsolete source; wrong effective-date filter; wrong version ranked higher.	Confirm authority and inspect retrieved versions.
A sensitive fact appears.	Excessive source permissions; wrong identity; filter failure; another data path.	Trace identity, access, retrieval, and output exposure.
A relevant passage yields a wrong answer.	Misinterpretation; missing exception; conflicting context.	Compare the answer with an owner-approved expected result.
The assistant refuses an answerable question.	Overrestriction; missing metadata; authorization mismatch; overly cautious response.	Determine whether the user's approved evidence was available.

Microsoft distinguishes evaluation of retrieved context from evaluation of the generated response. Its end-to-end guidance also treats groundedness and correctness separately. Eldris draws the following implication: an answer can follow its retrieved sources and still be unsuitable if the sources are outdated or inapplicable. [4, 5]

More connected data and more material inside a prompt are different interventions. Neither should be treated as a guarantee. Liu and colleagues demonstrated position-sensitive performance in long-context tasks using the models they studied. That historical result supports testing how a system uses evidence; it is not a current benchmark for a particular enterprise assistant. [7]

Diagnose before expanding

Trace the missing or incorrect answer to the stage that failed.

Source exists?	If no: fill the approved information gap.
User is eligible?	If no: verify intended access and enforcement.
Evidence retrieved?	If no: inspect indexing, filters, and ranking.
Answer appropriate?	If no: inspect applicability and interpretation.

Hold the task constant. Retest the diagnosed constraint after the change.

ELDRIS / ENTERPRISE AI ADOPTION GAP

© Eldris Inc

Figure 2. Investigate before expanding access. Eldris diagnostic framework. Check whether the required source exists, is eligible, is retrieved, and is correctly used before choosing an intervention.

A stronger model remains a legitimate option when answer generation is the constraint. Better indexing, ranking, or document structure can be appropriate when retrieval is the constraint. Source cleanup is appropriate when the governing evidence itself is wrong or ambiguous. Compare these alternatives on the same questions rather than credit the last change made to the system.

Choose the next deployment step

Eldris recommends three paths. Each has an evidence requirement and an explicit limit on what the organization is promising employees.

Launch with a curated collection

Choose this path when a maintained, approved collection supports a useful subset of employee questions and the user boundary can be enforced. Define the supported scope visibly, including when clarification or escalation is required. Evaluate the assistant on that scope and on nearby questions it should not answer definitively.

This path becomes weak when the promised service regularly needs excluded information. In that situation, a narrow launch can create avoidable refusals or incomplete guidance. Record those gaps rather than quietly implying full coverage.

Remediate the specific constraint

Choose this path when valuable questions cannot be handled reliably because source authority, access, or maintenance is unresolved. Prioritize the affected workflow and collection; an organization-wide cleanup is not automatically necessary.

Examples include correcting an oversized access group, identifying the owner of a policy, resolving conflicting effective dates, or fixing a connector's permission propagation. Each action should have a retest showing that the diagnosed problem changed. Discovery restrictions may buy time, but their limits should be documented. [2, 6]


Expand access when the additional information earns its place

Choose this path when the existing collection misses valuable questions and the additional sources have reviewed access, clear authority, and sustainable ownership. Compare the expanded collection with the current one using the same intended users and task mix.

Expansion should add useful answers without introducing unacceptable disclosure, wrong-policy answers, or maintenance obligations. It need not improve every question. It must improve the service sufficiently to justify its incremental burden within the enterprise's agreed limits.

Choose the next step from the evidence

Apply the decision to a defined workflow and collection.



Different collections in the same enterprise can require different paths.

Figure 3. The next step follows the diagnosed constraint. Eldris decision framework. These paths can apply to different collections within the same enterprise; they are not mutually exclusive organization-wide strategies.

Build an evaluation that can change the decision

The evaluation should compare genuine implementations of the candidate collections. Hold the model and response configuration constant for the initial collection comparison where feasible. Record unavoidable differences. Then evaluate retrieval or model changes separately so the enterprise can attribute improvements more credibly.

Start with employee questions collected from the intended workflow. Add deliberate cases covering outdated documents, conflicting policies, country-specific rules, legitimate access differences, missing context, and questions outside scope. Maintain a held-out set that was not used to tune the assistant.

Information owners should define acceptable answers and applicable sources. Security owners should define prohibited exposure. Employees should assess whether responses help complete the task. Resolve reviewer disagreement explicitly; agreement among model-based judges alone does not establish that a policy answer is appropriate.

Track the following outcomes separately:

- **Task coverage:** questions the approved collection can support, before evaluating retrieval.
- **Accepted outcomes:** correct, applicable answers or appropriate clarification and escalation, classified separately.
- **Unsupported confidence:** definitive answers when necessary evidence or context is missing.
- **Access failures:** prohibited material retrieved or exposed, including through citations and summaries.
- **Employee effort:** time to reach a usable outcome, including checking sources and resolving exceptions.
- **Change handling:** behavior after a policy update, permission removal, or ownership change.

An overall quality average can hide a wrong answer concentrated in one country or an exposure affecting a small group. Report denominators and results by material role and policy segment. Repeat selected questions to assess response variability. Microsoft notes that generation is nondeterministic and recommends continued evaluation as documents and user questions evolve. [5]

Define thresholds before comparing results. The report does not prescribe a universal acceptable error percentage. Owners must specify which failure classes prevent launch and which permit a restricted service with escalation. Observing no access failure in a finite test is not proof that no failure is possible.

A permission-change test should include the actual interval between changing access and observing the change across the assistant's retrieval, caching, and output paths. A configured setting is insufficient evidence of the operational boundary.

Account for the operating burden

A deployment decision needs evidence about the work required to sustain its performance. Eldris recommends comparing total workflow effort with the employee's existing route, such as search followed by clarification from a policy owner.

Measure employee checking and correction time, owner review time, access-remediation effort, content maintenance, and incident handling alongside licensing and infrastructure costs. Avoid counting a faster initial answer as a saving when the user subsequently spends more time validating it.

Separate startup work from recurring work. A one-time policy inventory and a continuing process for updating versions create different financial obligations. Allocate shared governance costs transparently rather than attributing an entire enterprise cleanup to one assistant or treating it as free.

For a collection expansion, ask three questions: which previously unresolved tasks became answerable; how much additional checking and maintenance did expansion require; and which new exposure or applicability failures appeared? Those observations determine whether to keep, modify, or reverse the expansion. This report contains no measured savings estimate.

A commissioned study could establish these quantities from real collections, role-based tests, task observations, and operating records. Interviews can explain organizational choices, but cannot substitute for direct evidence of access enforcement or answer behavior. An external participant panel can test comprehension of scope and clarification designs; it cannot establish the state of a sponsor's permissions or information.

The decision record to require

Before approving the next step, require a short decision record containing the following evidence:

- The intended population, supported questions, and visible limits of the service.
- The approved source collection, governing versions, accountable owners, and unresolved gaps.
- Verified identity and permission behavior, including changes and revocations.
- Comparison results, material segment differences, failure examples, and test limitations.
- The employee and owner effort required to operate the service.
- The selected path, reason, accountable approver, and conditions that trigger reconsideration.

A curated launch is defensible when it delivers a useful bounded service with maintained sources and tested boundaries. Remediation is defensible when it targets an identified constraint and leads to a meaningful retest. Broader access is defensible when additional information improves required outcomes and the enterprise can sustain its controls and maintenance.

The enterprise does not need every document to be perfect before obtaining any benefit from AI. It does need to know which information supports the promised service, who may use it, and how the assistant behaves when that evidence is missing or changes. **Approve the service that has been demonstrated for the intended users, then expand the promise as the evidence permits.**

Sources

Sources reviewed October 7, 2026. Product documentation describes vendor capabilities and responsibilities; it is not independent evidence of customer outcomes. Numbered references identify supporting sources. All figures and decision frameworks are Eldris analysis.

1. Microsoft Learn. [How data is protected and audited in Microsoft 365 and Microsoft Copilot](#). Access controls, information protection, and auditing architecture.
2. Microsoft Learn. [Configure a secure and governed foundation for Microsoft Copilot](#). Oversharing remediation, accountable ownership, and content hygiene.
3. Amazon Web Services. [Document-level access controls](#). Custom knowledge-base ACL filtering, authentication responsibilities, and permission limitations.
4. Microsoft Learn. [Retrieval-Augmented Generation evaluators](#). Separate retrieval and response evaluation dimensions.
5. Microsoft Azure Architecture Center. [Large language model end-to-end evaluation phase](#). Correctness, groundedness, variability, and continuing evaluation.
6. Microsoft Learn. [Restrict discovery of SharePoint sites and content](#). Discovery restrictions, permission distinctions, and scope limitations.
7. Liu, N. F., Lin, K., Hewitt, J., Paranjape, A., Bevilacqua, M., Petroni, F., and Liang, P. [Lost in the Middle: How Language Models Use Long Contexts](#). Research first submitted July 6, 2023; revised November 20, 2023. Historical evidence on context-position sensitivity in the tested models and tasks.