


FIELD GUIDE

Knowledge worth trusting.

A field guide to sources, access, and answers people can check.



A NOTE TO THE READER

A useful starting point.

An answer is only as useful as the information, permissions, and context behind it. A knowledge system needs a clear account of what it knows, who may use that knowledge, and what happens when the underlying material changes.

Use this guide before connecting a large document collection to an assistant. Start with a bounded set of recurring questions and the people responsible for the answers. The output should be an owned source collection, a defined audience, and a practical way to evaluate the resulting experience.

This is an operating framework, not a claim that any particular retrieval system guarantees accuracy or satisfies a legal requirement.

IN THIS EDITION

Begin with the questions.	3
Keep evidence attached.	4
Enforce access before the answer.	5
Plan for ordinary change.	6

KNOWLEDGE WORTH TRUSTING.

Begin with the questions.

A collection becomes useful when it serves a recognizable task.

Gather the questions people repeatedly need to answer in one area of the business. For each, identify the document or person that currently supplies the answer. Notice where the answer depends on a contract, region, effective date, or approval. Those conditions belong with the information.

Choose a collection with a clear owner who can resolve conflicting guidance. The smallest authoritative set is often a better starting point than an archive of everything ever discussed. Separate current guidance from drafts, historical records, and material that is useful only for a different audience.

Record genuine gaps. If an answer lives only in a person's memory, a search interface will not turn it into an approved policy. Decide whether that knowledge should be documented, whether the question should be escalated, or whether it falls outside the initial scope.

QUESTIONS TO TAKE INTO THE ROOM

Which recurring questions are worth answering first?

Who can say which source takes precedence?

What conditions change the meaning of an otherwise correct answer?

KNOWLEDGE WORTH TRUSTING.

Keep evidence attached.

Make a useful answer possible to inspect.

Preserve each source's title, location, relevant dates, owner, and conditions when preparing material for retrieval. An isolated excerpt can lose the qualification that makes it true. Upload time is not a substitute for an effective date, and a recent copy can still contain old guidance.

Design the answer so that a reader can connect important claims to supporting passages. A list of links at the bottom is not enough if it is unclear which source supports which statement. A citation should be checked during evaluation, including the exceptions and dates around the passage.

When sources are incomplete or contradictory, say what is missing and name a useful next step. The experience should distinguish documented guidance from interpretation. It should also be able to conclude that the approved collection does not answer the question.

QUESTIONS TO TAKE INTO THE ROOM

Can a reader trace each consequential statement to useful evidence?

Does the answer preserve the conditions in the source?

What should the experience show when evidence is missing or conflicting?

KNOWLEDGE WORTH TRUSTING.

Enforce access before the answer.

A restricted link does not protect a summary that already reveals its contents.

Define who is allowed to retrieve each source and enforce that boundary in the trusted application path. Hiding a document link after generating a response cannot undo information disclosed in the response. The retrieval and generation design must account for the user's actual authority.

Evaluate with different users asking the same question. Check both the returned source material and the final answer. Include combinations where a user can access a general policy but not a restricted appendix, and where the person's access has recently changed.

Keep the initial action boundary explicit. Explaining a procedure is different from carrying it out. If a later version can update a business record or send a message, that authority needs its own review, permissions, and recovery behavior.

QUESTIONS TO TAKE INTO THE ROOM

Where is the user's permission enforced?

Could an answer expose material the user cannot open?

How are changed or revoked permissions reflected in the system?

KNOWLEDGE WORTH TRUSTING.

Plan for ordinary change.

The information will move. The responsibility should remain clear.

Describe how edits, deletions, and access changes reach the knowledge system. Choose an update process that matches the freshness the workflow needs. If an update fails, an owner should be able to see what is stale and decide how the experience should respond.

Keep a recurring set of questions for checking changes. Include known answers, questions outside the collection, outdated guidance, conflicting documents, and material restricted to particular users. Assess relevance, support for claims, appropriate uncertainty, and access separately.

Treat corrections as operational evidence. A weak answer may reveal missing documentation, retrieval that lost context, or a policy nobody owns. Resolve the responsible layer. A more confident tone is not a repair for a source that was never available.

QUESTIONS TO TAKE INTO THE ROOM

Who owns source quality and failed updates?

What happens when a source is removed or superseded?

Which recurring checks tell us that the system still behaves as intended?

KEEP THE CONVERSATION USEFUL

The next considered step.

A dependable knowledge experience starts with a collection people own, permissions the system enforces, and answers that can be checked. Begin there, then expand what the evidence supports.

FURTHER READING

[Microsoft: Document-level access control in Azure AI Search](#)

A provider-specific reference for enforcing document access in search. Mechanisms and availability depend on the chosen platform.

<https://learn.microsoft.com/en-us/azure/search/search-document-level-access-overview>

[OpenAI: Evaluation best practices](#)

Further reading on evaluating the actual task and incorporating representative failure cases.

<https://developers.openai.com/api/docs/guides/evaluation-best-practices>

Written by Ayan. First edition, October 2026. This publication is general planning guidance. Project scope, fees, timing, and responsibilities are agreed separately in writing.

withayan.com / hello@withayan.com