



From Scattered Information to Shared Business Memory: Designing an AI-Ready Knowledge Foundation

A nontechnical architecture for turning scattered documents, decisions, procedures, and examples into governed company knowledge that people and approved AI systems can use.

ASSESS

Find the value

BUILD

Create the system

OPERATE

Keep it useful

In this field guide

Use this paper as a working reference for business decisions. The architecture is intentionally nontechnical: it names the goals, knowledge, responsibilities, controls, and operating practices that should exist around the technology.

01

Executive architecture

02

Four information layers that should not be confused

03

Why scattered information limits AI value

04

The six knowledge object types

05

The minimum knowledge record

06

Source authority and conflict resolution

07

Permissions belong inside the knowledge design

08

Assess: find the first knowledge domain

09

Build: create a minimum viable knowledge foundation

10

A practical example: project delivery knowledge

11

Operate: keep knowledge current and useful

12

Knowledge ownership model

13

Measurement framework

14

What not to put in shared business memory

15

Common design failures

16

The Assess, Build, Operate path

The central idea: company knowledge gives the system context, business goals give it direction, people keep consequential decisions, and managed operations keeps the whole system visible and improving.

This publication provides general business and design guidance. It is not legal, financial, clinical, employment, or security advice for a specific organization.

An AI-ready knowledge foundation is an organized set of approved company facts, procedures, decisions, examples, templates, and records with visible sources, owners, permissions, and review dates. It gives people and approved AI systems a shared place to find what the business currently believes, how work should be done, and where to go when the answer is uncertain.

It is not a folder dump, a copy of every document, or a chatbot trained on the company drive.

For this paper, an **AI agent** means software that can use an AI model to plan steps and work with connected tools. From here on, we use **specialized AI assistant** because the knowledge system should support bounded business roles with clear source and review rules.

The practical goal is **shared business memory**:

- People can find current guidance faster.
- Specialized AI assistants can prepare work from approved context.
- Answers preserve links to their sources.
- Conflicts and uncertainty route to named owners.
- Changes improve every approved use that depends on the knowledge.

This paper explains the architecture and operating model in nontechnical terms.

01 Executive architecture

The knowledge foundation has five connected stages.

Stage	Purpose	Business artifact
1. Source	Identify where authoritative information originates	Source register
2. Curate	Select, classify, clean, and resolve what is approved	Knowledge inventory and conflict log
3. Structure	Give information consistent meaning and ownership	Knowledge records, taxonomy, metadata
4. Use	Deliver sourced context to people and assistants	Answers, briefs, workflows, source links
5. Maintain	Review changes, feedback, access, and retirement	Review queue, change history, ownership cadence

The flow looks like this:

Source systems → approved knowledge → structured business memory → people and specialized AI assistants → sourced output → correction and maintenance

The arrow back from use to maintenance is essential. Questions reveal missing knowledge. Corrections reveal weak sources. Business changes create new versions. Without that feedback loop, a knowledge base becomes another stale location.

02 Four information layers that should not be confused

Companies often use "knowledge base" to mean several different things. A clear architecture separates four layers.

Layer	What it contains	Example	Management rule
Source system	Original business records and documents	Contract, CRM record, approved policy	Preserve authority and access controls
Knowledge foundation	Curated, structured guidance and references	Current onboarding procedure with owner and source	Maintain meaning, provenance, and currency
Operational context	Current facts for a specific case	This client's scope, open issue, and project status	Retrieve only for eligible work
Assistant working memory	Temporary context used while completing a task	Notes for the current readiness brief	Do not let temporary memory become policy

The distinctions protect accuracy. A conversational preference should not silently become company policy. A summarized procedure should not be treated as the approved source. A customer record should not be copied into a general company knowledge collection.

03 Why scattered information limits AI value

AI can generate fluent text without knowing which company source is current. When information is scattered, the system faces the same operating problems as the team:

- Several versions of a procedure appear valid.
- Important decisions live only in messages or memory.
- Templates have no clear owner.
- A current record is separated from the rule needed to interpret it.
- Access reflects folder history rather than business purpose.
- Corrections do not reach every tool or assistant.
- People cannot verify where an answer came from.

Adding a more capable model does not resolve source ownership. The business must decide what is authoritative, current, permitted, and appropriate for the job.

The [W3C PROV overview](#) defines provenance as information about the entities, activities, and people involved in producing data or another item. Provenance helps users assess quality, reliability, and trustworthiness. A small company does not need to implement a technical provenance standard to use the principle. It needs to preserve where information came from, who approved it, when it became effective, and what it supersedes.

04 The six knowledge object types

Instead of organizing everything only by department or folder, identify what kind of knowledge each item represents.

1. Definitions and business language

These records explain terms, status values, service names, customer categories, and other language that carries company-specific meaning.

Example: what "ready for kickoff" means and which conditions must be true.

2. Procedures and checklists

These records explain how repeatable work is performed, where judgment enters, what exceptions exist, and what completion looks like.

Example: the steps and review points for opening a new client project.

3. Policies and constraints

These records define approved behavior, prohibited behavior, thresholds, responsibilities, and escalation requirements.

Example: which types of customer credit require finance approval.

4. Examples and templates

These show what good work looks like in practice. They should be approved, representative, and labeled with context.

Example: a strong account readiness brief for a standard service relationship.

5. Decisions and exceptions

These explain why the business chose a rule, how a known exception was handled, and whether the decision should guide future work.

Example: a documented exception for projects with an outside compliance reviewer.

6. Operational records

These are current case-specific facts from systems of record. They may inform an answer or workflow, but they should not automatically become durable company guidance.

Example: a customer's current contract term or an open project issue.

This object model helps a specialized AI assistant distinguish "how we usually work" from "what is true in this case."

05 The minimum knowledge record

Each durable knowledge item should carry enough context for a person to judge whether it applies.

Field	Question it answers
Title	What is this?
Purpose	Why does it exist?
Knowledge type	Definition, procedure, policy, example, decision, or template?
Authoritative source	Where is the approved original?
Owner	Who may approve or change it?
Audience	Who is expected to use it?
Access class	Who may see it?
Effective date	When did it become current?
Review date	When should the owner check it again?
Supersedes	Which older item does this supersede?
Related work	Which processes and assistants depend on it?
Escalation	Who resolves ambiguity or conflict?

Not every note needs this treatment. Apply it to knowledge that drives repeated work, answers, or consequential decisions.

06 Source authority and conflict resolution

Conflicts are inevitable. Resolve them through a source hierarchy rather than letting the assistant choose the most convenient wording.

An example hierarchy might be:

1. Signed agreement or formally approved policy.
2. Approved system-of-record value.
3. Current procedure owned by the responsible team.
4. Approved decision record or exception.
5. Reviewed example or template.
6. Informal notes, messages, or personal memory.

The hierarchy is business-specific. The important part is to document it and route unresolved conflicts to an owner.

Use a conflict record with:

- The question or decision affected.
- The competing sources.
- The owner responsible for resolution.

- Temporary handling instructions.
- The approved decision and effective date.
- Sources or assistants that need updating.

A specialized AI assistant should be allowed to say that approved sources conflict. It should not hide uncertainty behind a smooth answer.

07 Permissions belong inside the knowledge design

Do not build a universal collection and try to add access rules later. Access is part of the meaning of company knowledge.

Use a simple classification model suited to the company:

Class	Example	Typical handling
Public	Published service description	Broad approved use
Internal	Standard operating procedure	Employees and approved systems with a business need
Confidential	Customer commercial information	Named roles and eligible workflows only
Restricted	Personnel, health, payment, legal, or credential data	Separate controls and qualified approval

For every specialized AI assistant, define:

- Which collections it may search.
- Which case records it may retrieve.
- Whether it can quote, summarize, or only reference a source.
- Which information may appear in the output.
- Who may receive that output.
- Which questions must route to a person.

The [NIST Privacy Framework](#) is a voluntary tool for identifying and managing privacy risk through enterprise risk management. Its practical relevance here is that information flow, purpose, roles, and outside providers should be understood before personal data is routed through an AI system.

08 Assess: find the first knowledge domain

Do not begin by organizing every file. Start with a business result that depends on repeated knowledge.

Good first domains have:

- Recurring questions or preparation work.
- A limited set of identifiable sources.
- A named process and knowledge owner.

- Existing examples of acceptable output.
- Clear cases that should escalate.
- A result people can review.

Examples include project delivery standards, service onboarding, product and inventory guidance, internal operating procedures, or account review preparation.

The knowledge discovery interview

Ask the people who perform and review the work:

1. Which questions repeat?
2. Which answers depend on company-specific context?
3. Where do you look first?
4. Which sources do you trust, and why?
5. Where do sources disagree?
6. What important guidance exists only in someone's memory?
7. Which information is sensitive or role-restricted?
8. Which answer would be harmful if it were stale or incomplete?
9. Who can approve a correction?
10. What should the system say when it cannot establish an answer?

Build a source register

Source	Owner	Authority	Access	Known issue	Used by
Approved procedure library	Operations lead	Current process guidance	Internal	Review dates incomplete	Onboarding and delivery roles
Customer system	Account owner	Current account facts	Confidential	Some fields inconsistently used	Account readiness role
Decision log	Leadership owner	Approved exceptions	Role-based	Older decisions lack expiry	Knowledge guidance role

The first assessment does not need to solve every issue. It should reveal whether a bounded, trustworthy knowledge set can be created.

The article on [building an AI-ready company knowledge base](#) provides a shorter implementation guide. An [AI Opportunity Audit](#) can place the knowledge work inside the larger business case and decide whether a knowledge foundation, workflow change, product configuration, or another path should come first.

09 Build: create a minimum viable knowledge foundation

Step 1: Select the first questions and outputs

Collect representative questions and work products from the chosen domain. Group them by intent, not only by wording. For example:

- Definition: "What does ready mean?"
- Procedure: "What happens after approval?"
- Exception: "What if a required item is missing?"
- Evidence: "Which source supports this answer?"
- Ownership: "Who decides this case?"

The question set becomes part of testing later.

Step 2: Curate the sources

For each source:

- Confirm the owner and authoritative version.
- Remove or label duplicates.
- Separate current guidance from archives.
- Identify restricted content.
- Record effective and review dates.
- Preserve the original location.
- Note unresolved conflicts.

Do not rewrite all source material into a new format unless the work needs it. The knowledge foundation can reference approved originals while adding the structure needed to find and apply them.

Step 3: Create the business taxonomy

A taxonomy is a consistent way to describe the knowledge. Keep it aligned with how the company talks about work.

Useful dimensions may include:

- Business area.
- Process stage.
- Customer or service type.
- Knowledge type.
- Risk or review level.
- Owner.

- Audience.
- Effective status.

Avoid a deeply technical or overly detailed structure. If employees cannot predict where something belongs, maintenance will fail.

Step 4: Define answer and evidence rules

Every knowledge-enabled assistant should follow rules such as:

- Use only approved sources for company-specific claims.
- Include source links or identifiers in the result.
- Prefer current authoritative sources over examples.
- State when sources conflict or do not answer the question.
- Do not infer policy from a single past case.
- Route sensitive, consequential, or ambiguous questions to named people.
- Distinguish company guidance from general information.

The representative article [Give the team one dependable answer](#) shows how those rules fit a practical workflow.

Step 5: Build the test set

Use four test categories:

Category	What to test
Answerable	Current source clearly supports the expected answer
Ambiguous	Several sources or interpretations could apply
Outdated	An older source conflicts with the current version
Prohibited	The user or assistant should not receive the information

For each test, record the question, eligible user, expected source, expected answer elements, required escalation, and unacceptable behavior.

Test source selection as carefully as writing quality. A polished answer from the wrong source is a failed result.

The [NIST Generative AI Profile](#) calls attention to provenance, data privacy, pre-deployment testing, and monitoring. The company knowledge implementation should make each of those visible in plain operating artifacts.

Step 6: Connect only approved uses

Start with a limited group of people, questions, and assistants. Keep the current manual path available. Require review where a wrong answer could create a customer, financial, access, legal, employment, safety, or clinical consequence.

A [specialized AI assistant implementation](#) should use the knowledge foundation through a defined role and eligible task, not receive unrestricted access because the connection is technically possible.

10 A practical example: project delivery knowledge

Suppose a service company wants consistent project starts and weekly readiness reviews.

Sources

- Signed scope and approved change records.
- Project delivery procedure.
- Service definitions and status meanings.
- Standard project template.
- Approved examples of readiness briefs.
- Current project and issue records.

Durable knowledge

- What each project status means.
- What must exist before kickoff.
- Who owns each exception.
- Which changes require customer or leadership approval.
- What a complete readiness brief contains.

Operational context

- This project's current scope.
- Current milestones and owners.
- Open issues and missing inputs.
- Recent approved changes.

Specialized assistants

Role	Knowledge use	Output
Intake preparation	Applies current kickoff requirements and templates	Prepared workspace and gap list
Readiness preparation	Combines current records with delivery standards	Weekly brief with source links
Exception coordination	Uses ownership and escalation rules	Routed questions and unresolved decisions

Human responsibilities

The delivery leader owns scope interpretation, staffing, tradeoffs, customer commitments, and final action.

Knowledge owners approve procedure changes. The managed operator reviews failures, access, corrections, and source freshness.

This is shared business memory in operation. The foundation supports several roles while preserving the difference between durable guidance and current case facts.

11 Operate: keep knowledge current and useful

Knowledge operations should have an explicit cadence.

Continuous signals

Capture:

- Questions the system could not answer.
- Answers reviewers corrected.
- Sources that conflicted.
- Requests denied by access rules.
- New terminology or exception patterns.
- Workflows affected by a changed source.

Weekly triage

The knowledge owner or delegate reviews high-impact corrections, unresolved questions, and urgent source changes. Assign each item an owner and due date.

Monthly maintenance

Review stale items, recurring questions, top correction categories, access changes, and upcoming review dates. Update the test set when new knowledge is approved.

Periodic portfolio review

Decide which domains should expand, remain stable, be redesigned, or be retired. Remove access and dependencies for knowledge no longer needed.

The [ISO/IEC 42001 standard](#) describes establishing, implementing, maintaining, and continually improving an AI management system. Knowledge maintenance is part of that larger operating system because company context directly affects system behavior.

12 Knowledge ownership model

Role	Responsibility
Business owner	Defines why the knowledge domain matters
Knowledge owner	Approves meaning, source, and changes
Source-system owner	Maintains authoritative records and access
Process owner	Explains how knowledge is applied in work
Reviewer	Judges answers and prepared outputs
Managed operator	Monitors feedback, tests, dependencies, and changes
Technical or privacy adviser	Reviews data flow, vendors, permissions, and retention

One person may fill multiple roles. Avoid the phrase "everyone owns knowledge." Shared use still requires named decision rights.

13 Measurement framework

Measure whether the knowledge helps the business, not how many documents were loaded.

Measure	Practical definition
Findability	Share of eligible questions that reach a relevant approved source
Grounding	Share of reviewed answers supported by the cited source
Currency	Share of active knowledge within its review period
Resolution	Share of conflicts and corrections closed by an owner
Review burden	Time required to judge and correct outputs
Reuse	Number of approved people or workflows using the same maintained source
Adoption	Eligible work actually using the knowledge route
Maintenance load	Time and cost required to keep the domain healthy

Set thresholds from the risk and business need. Do not invent a universal accuracy target. A low-impact internal directory question and a consequential policy question should not use the same acceptance standard.

14 What not to put in shared business memory

Exclude or tightly separate:

- Credentials, secret keys, and authentication material.
- Personal notes with no approved business purpose.

- Unverified claims and informal speculation.
- Outdated documents that are not clearly archived.
- Customer or employee information outside the defined need.
- Model-generated text that has not been reviewed as company guidance.
- Temporary task context that should expire.
- Licensed or third-party content the company is not authorized to reuse.

When uncertain, preserve the source in its controlled system and give the assistant only the minimum reference or retrieval path needed for the job.

15 Common design failures

Upload everything

Volume is mistaken for readiness. The result is duplicate, stale, sensitive, and conflicting content. Start with a bounded domain and approved sources.

Treat search as governance

Finding a passage does not establish that it is current, authoritative, or permitted. Add ownership, provenance, access, and review.

Copy the same guidance into every assistant

Corrections drift. Use shared maintained sources and role-specific retrieval rules.

Let conversation memory become policy

A preference or past answer silently becomes durable guidance. Separate temporary working context from approved knowledge.

Hide uncertainty

The system produces a confident answer even when sources disagree. Require a visible conflict and escalation path.

Build without maintenance

No one reviews sources after launch. Assign operating ownership and measure currency, corrections, and unresolved questions.

Count documents instead of outcomes

The library grows while repeated work remains unchanged. Tie the knowledge domain to a business result and the specialized roles that use it.

16 The Assess, Build, Operate path

Assess

- Choose one business result and knowledge domain.
- Interview users, reviewers, and source owners.
- Inventory authoritative, conflicting, and restricted sources.
- Define representative questions and expected outputs.
- Decide whether the knowledge work is worth doing now.

Build

- Curate the first approved source set.
- Create the taxonomy and minimum knowledge records.
- Define source, answer, access, and escalation rules.
- Build a representative test set.
- Connect a limited group of people and specialized assistants.
- Run beside the current process before broader use.

Operate

- Review unanswered, corrected, and disputed outputs.
- Maintain owners, effective dates, access, and dependencies.
- Retest when knowledge, tools, or roles change.
- Measure business use, quality, currency, and maintenance cost.
- Expand, restrict, redesign, or retire deliberately.

The knowledge foundation is one layer of the broader [Managed AI Operations Blueprint](#). It gives the system context. Business goals give it direction. Human ownership and managed operation keep it trustworthy and useful.

Northern Logic can assess the first knowledge domain, build the foundation, and manage it as the company changes. Explore the [Business Knowledge Base service](#) or [book a free 15-minute AI assessment](#) with one repeated question, preparation task, or source problem.

17 Sources

Primary and authoritative references used throughout this field guide.

- [W3C PROV Overview](#)
- [W3C PROV Model Primer](#)
- [NIST AI Risk Management Framework](#)
- [NIST AI RMF Core](#)
- [NIST Generative AI Profile](#)
- [NIST Privacy Framework](#)
- [ISO/IEC 42001:2023 AI Management Systems](#)
- [CISA and NCSC Guidelines for Secure AI System Development](#)
- [OWASP LLM06:2025 Excessive Agency](#)