



The Assessment – 30 gate questions

Answer YES only on evidence; ties break downward. A dimension's level = consecutive YES count from its first question. Your organization's level = your lowest dimension.

Dimension 1 – Capture (how knowledge gets created)

- [] **C1.** A knowledge base exists, and the people who solve problems can add to it.
- [] **C2.** Capture happens *in the workflow*: articles are created or linked as part of working a case, under written standards — not "when someone gets a chance."
- [] **C3.** AI drafts candidate articles from resolved-case material; solver effort per capture is under a minute; humans curate rather than author.
- [] **C4.** Automated signals also fire capture — recurrence detection or gap reports — arriving with demand evidence attached.
- [] **C5.** For standard cases, capture-through-publish runs without a human initiating it; humans handle exceptions and policy.

Dimension 2 – Curation (how content gets judged)

- [] **U1.** New content is reviewed by someone before the organization treats it as trusted.
- [] **U2.** Review follows defined article standards, and reviewers are coached against them.
- [] **U3.** A formal gate exists: every candidate gets a disposition with a reason code, edit-distance is logged, and a disposition SLA is held.
- [] **U4.** Changes to *existing* articles pass a before/after DIFF review; retirement is gated, and archived content is never deleted.
- [] **U5.** Automated pre-checks (quality, dedupe, policy screens) triage the queue so humans only see what machines can't clear.

Dimension 3 – System of Record (structure and trust fields)

- [] **S1.** Articles live in one designated system — not scattered across drives, wikis, and chat pins.
- [] **S2.** Articles follow consistent, typed templates with fixed section structure.
- [] **S3.** Every article carries machine-readable governance metadata: stable ID, type, owner, lifecycle status, and applicability (product/version/domain/audience).
- [] **S4.** Content is versioned with retrievable history; lifecycle states are *enforced* — only validated content reaches any delivery surface.
- [] **S5.** Relationships between articles (supersedes, related, depends-on) are explicit and queryable — you can answer "what references the thing we just changed?" from structure, not memory.

Dimension 4 – Delivery (who and what consumes)

- [] **D1.** Humans can search the knowledge base and find what exists.
- [] **D2.** Agents use articles in-workflow, and you can measure reuse (link rate on resolved cases).
- [] **D3.** At least one machine consumes your knowledge — agent-assist or a customer-facing bot — drawing on the same corpus humans use, not a separate copy.
- [] **D4.** Machine answers cite article IDs, filter by applicability and validated-only status, and *refuse* below a confidence floor rather than improvise.
- [] **D5.** External AI agents (your customers') can consume governed, cited answers through an API, under the same guarantees.

Dimension 5 – Telemetry (what you can see)

- [] **T1.** Basic usage data exists (views, searches) and someone looks at it.
- [] **T2.** Participation and reuse rates are tracked against program standards.
- [] **T3.** Answer events are attributed per channel with unit costs — cost per answer is computable — and curator edit-distance is logged.
- [] **T4.** Unanswered queries are logged and clustered into gap reports; blended cost per answer and the shift-left ratio are on a dashboard that is *sent* to decision-makers.
- [] **T5.** Telemetry drives automation: drift detection and gap-to-capture requests fire without a human noticing first.

Dimension 6 – Governance (rules and ownership)

- [] **G1.** A named person is responsible for the knowledge base as a whole.
- [] **G2.** Content health work (reviews, cleanup) happens on a cadence, not just when things break.
- [] **G3.** Every article has an accountable owner, and ratified content standards exist — including the rule that nothing AI-generated enters the record ungated.
- [] **G4.** A lifecycle policy is enforced with signal-based review, and governance decisions land in an append-only log.
- [] **G5.** Thresholds and policies are tuned quarterly *from the log* — governance evolves on records, not opinions.

Scoring grid

Dimension	Consecutive YES from Q1	Level
Capture		
Curation		

Dimension	Consecutive YES from Q1	Level
System of Record		
Delivery		
Telemetry		
Governance		
Organization level (lowest row)		
Weakest dimension(s)		

Assessed by: _____ **Date:** _____ *(Keep every completed grid; never edit a past one.)*

→ Now read your level in `levels-and-next-moves.md` , then fill `roadmap-one-pager.md` .

Contextkeeping is a KCS-aligned operating model. KCS® is a service mark of the Consortium for Service Innovation; this material is not affiliated with or endorsed by the Consortium.