



The Contextkeeping Maturity Model

Where is your knowledge program on the road from institutional knowledge to agentic AI?

Most support organizations can't answer that question — which is exactly why "add AI to the knowledge base" projects fail. AI amplifies whatever knowledge operation you already have. If yours is a graveyard of stale articles, AI gives you confident, cited, *wrong* answers at scale. Find your level below, then look one row down: that's your next move.

Level 0 - Institutional. Knowledge lives in heads, closed tickets, and Slack scrollback. Every hard problem gets re-solved from scratch; when your best agent leaves, their knowledge leaves with them. *You're here if:* new hires learn by shadowing, and "check with Dana" is a documented process.

Level 1 — Documented. A knowledge base exists. Writing articles is "when you get a chance," which means never during a busy queue. Content decays silently; nobody trusts it enough to check it first. *You're here if:* your KB's most-viewed article was last updated two product versions ago. *Key metric to start watching:* % of resolved tickets with a covering article.

Level 2 — Process-driven. Classical KCS or equivalent: capture in the workflow, article standards, coaching, content health reviews. This is real maturity — and it's where most certified programs are stuck, because it assumes humans write every article and humans read every article. Both assumptions just expired. *You're here if:* your program works but scales linearly with agent effort. *Key metric:* participation and reuse rates.

Level 3 — AI-assisted. AI drafts articles automatically from resolved-case payloads; humans stop authoring and start *curating* — approving, editing, or rejecting candidates through a gate in minutes, not hours. You track how much reviewers change AI drafts (edit-distance), so drafting quality measurably improves. *You're here if:* capture takes an agent under a minute and publish happens same-day. *Key metrics:* capture-to-publish latency; edit-distance trend.

Level 4 — Governed. A true system of record: every article versioned, owned, tagged with product/version/audience, and carrying a review date. Updates to existing knowledge pass a before/after review gate. Knowledge is delivered to machines as well as humans — agent-assist and chatbots answer *with citations*, refuse when coverage is missing, and unanswered questions automatically become capture requests. The loop closes itself. *You're here if:* you can answer "what's stale, what's missing, and what did AI answers cite last week?" from a dashboard. *Key metrics:* blended cost per answer across all channels; % of corpus within review date; gap-closure latency.

Level 5 — Agentic. Orchestrated AI workflows run capture-through-publish end to end for standard cases; humans govern exceptions, policy, and quality thresholds. Your knowledge base is infrastructure — consumed

by your own AI agents and, increasingly, your customers' agents. *You're here if:* honestly, almost nobody is yet. That's the point.

The uncomfortable rule of thumb

You cannot skip levels. Bolting Level 4 tooling onto Level 1 content produces expensive hallucinations with your logo on them. The good news: Levels 1→3 can be crossed in weeks with the right operating model — no enterprise CMS required.

Find your level, then your next move

→ **The Maturity Self-Assessment** (contextkeeping.com/toolkit) — the 30-question scored instrument that places you honestly, dimension by dimension, with your next moves per level. → **The Contextkeeping Toolkit** (contextkeeping.com/toolkit) — the templates, curation-gate checklists, and measurement workbook to start climbing today.

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