



Levels & next moves

Find your organization level (your lowest dimension), read your page, execute the moves. Each move names the module that executes it. You cannot skip levels — the moves climb one rung at a time, on purpose.

Level 0 — Institutional

You're here if: knowledge lives in heads, closed tickets, and chat scrollbar; new hires learn by shadowing; "check with the SME" is a documented process. **Failure modes:** every hard problem re-solved from scratch; your best agent's resignation is a knowledge-loss event; onboarding takes quarters. **Watch:** % of resolved cases with any covering article (it's near zero — write it down anyway; it's your before-picture). **Next moves:** designate the one system of record and stop scattering (Module 2 §8 adapters) · stand up the five typed templates (Module 2) · name an accountable KB owner (Module 5, ownership policy) · capture the Baseline Snapshot (Module 6) — even mostly-empty, it starts the clock.

Level 1 — Documented

You're here if: a KB exists; writing is voluntary and after-the-fact; the most-viewed article was last updated two product versions ago; nobody checks the KB first. **Failure modes:** silent decay; duplicate and contradictory articles; trust so low that agents answer from memory *around* the KB. **Watch:** % of resolved cases with a covering article · % of corpus with an owner · article age distribution. **Next moves:** type the corpus — run the metadata-extraction prompt (Module 3) over existing articles and expect a fifth (20%) to be retire-candidates; that's the migration working · ratify the content standards and lifecycle policies (Module 5) · install ownership with an orphan sweep (Module 5) · begin the review-date ledger (Modules 5 + 6).

Level 2 — Process-driven

You're here if: classical KCS or equivalent — capture in the workflow, standards, coaching, health reviews. Real maturity; certified, probably. And stuck: the program scales linearly with human effort, and both of its load-bearing assumptions (humans write, humans read) have expired. **Failure modes:** the capture tax quietly re-wins — throughput pressure beats article quality; your best people carry the program; leadership asks "why do we pay for this" every budget cycle. **Watch:** participation and reuse rates · capture-to-publish latency (in days at this level; the next level makes it hours). **Next moves:** wire AI drafting to your resolved cases and run the evaluation protocol before trusting it (Module 3) · convert reviewers into gatekeepers — dispositions, reason codes, edit-distance (Module 4) · reframe your measurement to cost per answer *before* leadership asks again (Module 6) · your existing standards map onto Module 2 in an afternoon — keep what works.

Level 3 — AI-assisted

You're here if: AI drafts from case payloads; solver effort is under a minute; humans curate through a gate in minutes; same-day publish is normal. **Failure modes:** gate becomes a bottleneck (queue silts up) or a rubber stamp (quality leaks) · candidates flow but existing content rots — capture is modern, curation of the *old* corpus isn't · drafting quality drifts unnoticed on a model change. **Watch:** capture-to-publish latency · edit-distance trend · candidate acceptance rate · queue age. **Next moves:** enforce lifecycle end-to-end — validated-only delivery, DIFF-gated updates, gated retirement (Modules 5 + 4) · instrument gaps: log unanswered queries, cluster them, feed capture (Modules 3 + 6) · put blended cost per answer and shift-left on a dashboard someone *receives* (Module 6) · send the Inversion Brief before assisted metrics "worsen" (Module 6).

Level 4 — Governed

You're here if: a true system of record — versioned, owned, lifecycle-enforced; machine delivery answers with citations and refuses honestly; unanswered questions automatically become capture requests; you can answer "what's stale, what's missing, what did AI answers cite last week?" from a dashboard. **Failure modes:** complacency — the loop runs but nobody tunes thresholds, and override patterns go unread · governance theater: logs exist, reviews don't · the sixth channel (customers' AI agents) arrives before you've priced or governed it. **Watch:** blended cost per answer · % corpus within review date · gap-closure latency · citation freshness · override rate. **Next moves:** tune quarterly from the records layer — reject-reason mix, override patterns, gate throughput (Modules 4 + 5) · pilot automated capture-through-publish for your narrowest, most deterministic case class, with humans on exceptions (toward Level 5) · open the API/MCP surface deliberately: applicability filters, citations, confidence floor, per-answer cost tracked from day one (Module 6, sixth channel) · re-run this assessment — your weakest dimension has moved.

Level 5 — Agentic

You're here if: honestly, almost nobody is — orchestrated workflows run capture-through-publish for standard cases end to end; humans govern exceptions, policy, and thresholds; your knowledge serves your own AI agents and your customers' as governed infrastructure. **Failure modes:** automation outrunning its gates — the temptation to remove the human from *update* flows, where the Precision Paradox bites hardest · treating the curation dividend as a headcount cut, which stalls the flywheel that got you here. **Watch:** exception rate (share of flow needing a human — falling is maturity, zero is a warning) · refusal correctness · cost per governed answer · external-agent answer share. **Next moves:** hold the line — the boring core stays boring, gates stay on writes, and the AI layer stays swappable · publish your practice: at this level your knowledge operation is itself a competitive story.

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.