



Content Standards Policy — [ORGANIZATION]

Fill the brackets, strike what doesn't apply, ratify, and date it. Target length after filling: two pages.
Effective date: [DATE] · Review: quarterly · Owner: [SUPPORT LEADER].

1. Purpose and scope

This policy governs what may enter [ORGANIZATION]'s knowledge system of record, in what form, and by whose hand. It applies to all knowledge content for [DOMAINS IN SCOPE — e.g. "product support and internal IT"], across every authoring source: human, AI-drafted, or migrated.

2. Typed content only

Every article is exactly one of the approved types — **issue/resolution, how-to, FAQ, known error** — using the corresponding template and section structure (Module 2). Content that fits no type either isn't knowledge (don't add it) or reveals a missing type (amend this policy deliberately).

3. The probabilistic-write rule

No AI system writes to the system of record. AI output enters only as a *candidate* in the queue and becomes an article only through the new-candidate gate; AI-proposed changes to existing articles pass the update-DIFF gate. This rule has no exceptions and survives every tooling change. (*Human writes bypass nothing either: the same gates apply — the gate is about protecting the record, not about distrusting machines.*)

4. Language and quality

- Articles preserve the **customer's language** — symptoms and questions as reported, including exact error text. Editing for clarity must not sanitize findability away.
- Articles are **sufficient to solve** — and no longer. One issue, task, question, or defect per article.
- Sections are **self-contained** (they will be retrieved in isolation) and use their template's headers verbatim.
- Prohibited in all content: personal data (names, emails, account identifiers), customer-identifying context, credentials, internal-only URLs in customer-audience articles, and claims not verified against a source. [ADD ORG-SPECIFIC PROHIBITIONS — regulated-industry language, legal-required phrasing, etc.]

5. Metadata is mandatory

An article without complete frontmatter (per `metadata-schema.yaml`) does not get validated — the schema check is part of the gate, not a suggestion. In particular: `owner` never empty, applicability (`product` , `product_version` , `domain` , `audience`) always structured, `volatility` always set at creation.

6. Delivery rule

Only `validated` content may ground automated answers, appear in agent-assist, or display as trusted knowledge. Draft and archived content are excluded from every delivery surface. Answer systems must cite article ids and must refuse rather than improvise when coverage is missing. [NAME THE SURFACES: help center, chatbot, agent-assist panel, API.]

7. Exceptions

Deviations from this policy are permitted only as **logged overrides** in the decision log, with actor and reason. Three identical overrides in a quarter = this policy is wrong; amend it at the quarterly review rather than continuing to override.

Ratified by: [NAMES] · [DATE] · Governance Pack, Module 5 · Contextkeeping Toolkit

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.