



Content Standards Policy – Meridian Systems

The toolkit's worked example: the template filled in, signed, and in force

This is the toolkit's worked example: the `content-standards-policy.md` template, filled and ratified for the fictional Meridian Systems (the same company whose articles fill Module 2's `examples/`). Every bracket resolved, nothing struck that wasn't decided out loud. Use it to see what "done" looks like — then fill your own. Effective date: 2026-04-30 · Review: quarterly · Owner: T. Okafor, Support Director.

1. Purpose and scope

This policy governs what may enter Meridian Systems' knowledge system of record, in what form, and by whose hand. It applies to all knowledge content for **customer product support and internal IT**, across every authoring source: human, AI-drafted, or migrated. Out of scope, deliberately: engineering design docs and sales collateral, which have their own owners and their own rules.

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). *Decision noted at ratification: release notes are not a fifth type; they feed the how-to synthesis prompt instead.*

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 — including the Zendesk answer bot, which reads `validated` content and writes nothing. *(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.

Meridian-specific: no customer company names even as "anonymized" anecdotes; no screenshots containing tenant identifiers; security-hardening steps for the endpoint agent are `audience: internal` only and never appear in customer-facing articles.

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. Meridian's domain taxonomy: **endpoint-support · billing · platform-admin · product-defects**, owned respectively by **team-endpoint · team-billing · team-platform · team-defects** (per the ownership policy's table).

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. **Meridian's surfaces:** the help center, the Zendesk agent-assist panel, the customer-facing answer bot, and the partner API — all four read the same corpus under the same filter.

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. *First quarterly review scheduled 2026-07-30; the review reads the log, not opinions.*

Ratified by: T. Okafor (Support Director) · J. Rivera (Knowledge Lead) · 2026-04-30 · 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.