



Curator Training — 90-minute session outline

For the 2–3 people taking the curator hat. Prerequisites: they've read Module 4's guide; you've printed the three gate checklists and pulled Module 2's example articles.

0:00–0:10 — Why this role exists

The one rule (nothing probabilistic writes to the record ungated) and why it's architecture, not ceremony: the Precision Paradox — the better the answers sound, the more the worst one costs. Their job in one line: *you decide what becomes true here*. Time budgets on the wall: clean approve ≤ 3 min, edit-approve ≤ 10, reject ≤ 2.

0:10–0:30 — Anatomy of a good article (using Module 2's examples)

Walk the gold-standard issue-resolution example ([K0-000142](#)): frontmatter as governance, sections as the payload, customer language, self-containment, negative space. Then the ten anti-patterns (Module 2 §7) at speed — curators reject by pattern recognition, not by rulebook lookup.

0:30–1:00 — Live gate practice (the heart of the session)

Six candidates, prepared in advance — the trainer's real work is building this set:

1. A clean approve (well-grounded, complete) → **A1** in under 3 minutes.
2. A light-edit case (support-speak crept in; one `[NEEDS-VERIFICATION]` to resolve) → **A2**, log edit-distance, feel the 0–3 scale.
3. A duplicate of an existing article → **R2 with the KO- id cited** — practice actually searching the corpus first.
4. A plausible, fluent draft with a *wrong resolution step* planted → the session's money moment: fluency is not accuracy. **R3**, or escalate to the owner if they can't verify — both are right answers; waving it through is the only wrong one.
5. A PII leak (customer name + account id in the Issue section) → **R4**, and why "fix at publish" is banned.
6. An update-DIFF with a collateral edit hiding in an untouched section → reject the collateral, accept the intended change; meaning-change test discussed.

Each candidate: individual disposition first, then compare, then the checklist retrace. Disagreements are the curriculum — talk them out.

1:00–1:15 – Logging and the loop

One disposition each into the decision log, live. Where the data goes: edit-distance → drafting quality trend; reason codes → prompt and threshold tuning; their SLA → the dashboard. The point to land: *logging isn't admin overhead, it's the program's sensory system, and they are its nerve endings.*

1:15–1:30 – Operating rhythm + close

Daily 20–30 min batch beats drip-checking · queue age is their number · escalation path (owner verifies facts, curator still disposes, 5-day cap) · when to say "raise the capture threshold" (Module 4's anti-bottleneck rules — they are allowed, expected, to push back on volume). Close: first real shift is tomorrow, paired; solo by week's end; 30-minute retro after each of the first two weeks.

Refresh: rerun candidates 4–6 (fresh examples) for 30 minutes whenever a new curator joins or the prompts version-bumps.

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.