



# The Content Graduation Model

*One page: the ladder, the criteria, the pass, and the pilot — [machinereadyknowledge.com/graduation](https://machinereadyknowledge.com/graduation)*

---

**The trade.** Stop managing content as inventory in channels; manage it as flow between them. Every channel has an owner — this model is the missing layer that owns the *movement* of knowledge: promotion when demand and quality earn it, enrichment when usage calls for it, archival when the job is done.

**The ladder.** Conversation → validated answer → knowledge article → official documentation → training & rich formats → **the product itself** (in-product help, better defaults, fixes — knowledge with zero reading required). Movement goes down as well as up: superseded, invalidated, or unnecessary content retires with a redirect. A shrinking workaround library is a quality metric.

**The criteria — four signal families.** **Demand** (accesses, accepted answers, reuse, case-attach, recurrence — **MTTE**, mean time to event: how quickly the same issue resurfaces). **Quality** (validated truth, completeness against the destination's standard). **Freshness** (activity, trend, a re-evaluation horizon on every asset — nothing is permanently finished). **Outcome** (answers delivered by content instead of the assisted queue; the standing product question: *should this content need to exist?*). Instrument them as a small, boring record on every candidate — criteria in metadata scale; criteria in opinions don't.

**The pass — machines detect, humans decide, machines carry.** Harvest on a schedule → evaluate against ratified thresholds → **validate at a human gate** (non-negotiable — the difference between a knowledge layer and a rumor mill) → rewrite to the destination's standard → publish → **close the loop**: post the graduated link where the question lives. Closure is what teaches contributors the ladder exists. The same pass runs retirement — with the same gate.

**Why AI raises the stakes.** Retrieval answers from the corpus it's pointed at. The graduated corpus is trust-tiered (retrieval weights by earned tier; archived content is never served), provenance-carrying, and continuously pruned — the version of your knowledge that is safe to ground a confident answer engine in. Every question the assistant can't answer well re-enters the loop as a demand signal.

**Start embarrassingly small — the 1 / 5 / 2 pilot.** One flow (forum → knowledge base, usually), five candidates per pass, two passes, entirely by hand, every threshold tested against real content, every loop closed publicly. Then instrument, automate the carrying, add the second flow, and stand up the content × product cadence.

*The full machinery — the model paper, eleven-flow catalog, trigger-criteria worksheet, graduation-record schema, pass runbook, gate line-items, evolve cadence, and champion deck — is the Content Graduation Pack at [contextkeeping.com/toolkit](https://contextkeeping.com/toolkit).*