

Committee quorum #6 — retrospective

vade-coo

2026-04-26

Table of contents

What happened	1
What worked	3
What surprised me	4
What broke and got fixed mid-flight	5
What the pattern is — and isn't	6
Open questions	7
Forward-looking	8

Authored 2026-04-26 by COO (commissioner instance for the arc; cloud session run-2026-04-26T101413, continuation of run-2026-04-26T122318). Scope: the full arc of committee quorum #6 — repo organization sweep — from issue #174 commissioning through 11-instance convergence to post-merge in PR #177 + MEMO-2026-04-26-09. Sibling retros: quorum #1 (2026-04-23_committee-quorum-1.md, protocol designing itself) and quorum #3 (2026-04-24_committee-quorum-3.md, memo-system-transition design-decision arc).

What happened

Issue #174 commissioned a quorum to deliberate a repo-reorganization sweep covering `vade-coo-memory` and `vade-runtime` to fix three problems: kind-mixing in folders, severely stale front-door docs (`README.md` + `context/bootstrap_plan.md`), and post-shipment artifact accumulation in `coo/_drafts/`. The commissioning issue carried a comprehensive read-only audit (K1–K28 file-kind taxonomy, per-folder findings, 5-round suggestion, 7 deliberation seams) as a starting point; the committee was free to restructure.

Arc shape (11 instances + commissioner + 2 adversarial auditors, ~3.5 hours calendar on 2026-04-26):

- **Instance #1** (commit `666b3c5`, +816/–70): full-context, no sub-agents. Filled the empty scaffold; declined the K1–K28 taxonomy as a durable artifact (reasoning per-folder is lighter); inverted the audit's round ordering (front-door first, archive-sweep second, on a daily-read-cost argument); recommended `bootstrap_plan.md` archive option (a) over (b) and (c); introduced `coo/_evidence/` for K10 sub-agent reports; threshold-rule for `coo/retrospectives/` rather than prophylactic split; new `coo/draft_lifecycle.md` short SOP. 7 deliberation seams added to the audit's 7. Substantive round 1/10.
- **Instance #2** (commit `0f07440`, +202/–62): three substantive catches — Phase 4 meta-skill cited in MEMO-2026-04-25-01 (promote, not archive); `docs/auto-tag-routine.md` cited in 4 memos as a living append-log (keep, don't collapse); three numeric reconciliations. Introduced the §“Memo citation safety” section with grep-before-relocating methodology and three drift-handling rules. Round 2/10.

- **Instance #3** (commit 2c53eed): `adoption_research.md` cited in MEMO-2026-04-26-07 §Decision (instance #2 missed it under the same false-dichotomy class as Phase 4 — third-for-three on audit-named promotion targets); three numeric errors in the citation-safety table (bare-basename vs strict-path grep conflation); seven `_drafts/` archive targets missing from the table. Survey scope clarification added (memo-body cites vs non-memo durable-doc cites). Round 3/10.
- **Instance #4** (commit 5277d0f): two methodology refinements on path-update counts (line-anchor `<basename>:<line>` refs and tracker-table data rows are basename-keyed, unaffected by directory-only relocation; instance #3 over-counted both); 1-file archive sub-bucket folded into `shipped-stagings/`; net file-count arithmetic fix. Codified the flat-vs-bucket rule. Round 4/10.
- **Instance #5** (commit 69c9a08): three substantive catches — Round 2 missing inline self-references in `adoption_research.md` (5 edits at lines 128–132); Round 2 missing inline path-string updates in project-history retro (3 line-cites); **Round 3 K10 move conflicted with MEMO-2026-04-23-01 essay voice-preservation** (essay opening names K10 files by date-prefixed basename; protocol-binding catch surfaced as taste seam for 4 rounds). Methodology extended to 4 reference shapes; new disposition (iv) basename preservation. Round 5/10.
- **Instance #6** (commit d2a46c6): **second protocol-binding catch on the same Open Seam #5 thread instance #5 thought resolved**. Companion transcript `coo/foundations/2026-04-22_transcript.md` references the K10 files at full path in three turns (Ven’s user-turn co-author briefs at lines 826–827 and 974–977; COO’s assistant-turn report at lines 740–741). Transcript verbatim invariant blocks both basename normalization AND inline edits; resolution is stay-in-place for the 2026-04-22 K10 batch. Methodology extended to 6 reference shapes (added (e) live-tooling refs, (f) intra-bucket batch refs); 7 dispositions (added (vi) bucket-level path-drift note, (vii) stay-in-place). Cross-repo grep claimed empirically clean. Round 6/10.
- **Instance #7** (commit e02bca8): two substantive catches. First, 5 inline path-string refs in `.claude/agents/` that instance #6’s shape-(e) extension missed (#6 covered `.claude/skills/**` and `.claude/commands/**` but not `.claude/agents/**`). Second, **cross-repo verification claim correction** — instance #6’s “zero hits” was wrong; `vade-runtime/scripts/integrity-check.sh:333` regex needs `_evidence/` added (load-bearing — without it Round 2 commits would false-positive F1 invariant on next integrity probe); `vade-agent-logs/sessions/2026/04/.../*.md` has 13 path-string references under disposition (iii). Surfaced the “verification-overclaim failure mode” as a process observation. Round 7/10.
- **Instance #8** (commit 7ddd4bc): one substantive catch — 2 `vade-runtime` live-tooling emit lines instance #7’s grep missed (`discussions-digest.sh:112` fires on every `SessionStart`-hook; `session-lifecycle.sh:151` fires on end-of-session summary; both emit `vade-coo-memory/context/agent-boot-discussions-check.md`). **Codified the verification-claim discipline** as a methodology rule (state the search-target list and the corpus searched explicitly). New OQ #9 added (Phase-3 specialist agent retention, surfaced by #7). Round 8/10.
- **Instance #9** (no commit): **first approval**. Independent exhaustive cross-repo verification using a methodologically distinct scoping (full-path Explore subagent — first sub-agent dispatch in the entire arc). Confirmed instance #8’s three carve-outs at exact line numbers; firmed up instance #7’s `vade-agent-logs` count of 13 (instance #8’s “sampled” characterization was itself an overclaim — the count was exhaustive). Three non-gating drift-flag observations recorded. Convergence to 1/2.
- **Instance #10** (no commit): **second approval, 2/2 quorum reached**. Independent re-verification using basename-only grep across all 4 sibling repos (a methodologically distinct scoping from #9’s full-path Explore). Same 3 carve-outs; no fourth class of cross-repo hits surfaced. Convergence to 2/2.

- **Instance #11** (no commit; **supernumerary**): the Quorum member routine re-fired `pull_request.assigned` for `@vade-coo` after #10's quorum — same trigger-prompt vs spawn-guide divergence pattern as quorum-#3 #9 (PR #91 was supposed to close it). Instance #11 handled correctly per case-law (treated as one more independent verification, didn't reset approvals). Used within-repo `.claude/**` archive-target enumeration as a third methodologically-distinct verification. Caught two non-gating observations: `.claude/commands/_lib/memo-index.sh:7` comment citing `coo/memo_system_transition.md` (within-repo equivalent of the verification-overclaim pattern); `_evidence/2026-04-24_project-history-retrospective/` dirname doesn't strictly satisfy the parent-artifact-basename rule. Declared 3/3 over-quorum.

Commissioner ratification: identity confirmed; compiled drift-flags + minimal-fix rewrites; presented three options to BDFL. **Ven commissioned adversarial audits** before ratifying (non-blocking); both passed: - safety-auditor: 8 findings, **0 BLOCK, 2 PASS WITH NOTE, 6 PASS** across MEMO -22-01 / -22-08 / -22-10 / -22-14 / -22-19 + F1 invariant + verbatim-transcript invariant. - emancipatory-auditor: **2/2 across all major sections**; 2 conditional 2/1s on `draft_lifecycle.md` SOP and citation- safety methodology that recover to 2/2 if Round 5 absorbs OQ#8.

BDFL ratified with option (c). Post-merge sequence executed: PR #175 squash-merged at `7d1770f`; post-merge PR #177 squash-merged at `6fabde6` with rename + 5 drift-flag fixes + MEMO-2026-04-26-09; commissioning issue #174 closed with citations; 4 follow-up issues filed (#178 OQ #3 `architecture_as_built`; #179 OQ #6+#8 bundled into `draft_lifecycle.md`; #180 OQ #9 Phase-3 specialist agent retention; #181 process — Quorum member routine reassigned-after-quorum bug).

Totals: 8 substantive rounds used (BDFL-bumped cap of 10); 3 approval rounds; 0 BDFL returns; 0 BDFL clarifications in-flight; 2 sub-agent dispatches (instances #9 + #10, both Explore at the convergence-test moment); 5 drift-flags absorbed at merge; 4 follow-up issues filed; 1 §8 summary memo; 1 retrospective.

What worked

The narrowing-catch signal reached zero across three consecutive rounds. Quorum #3's retrospective proposed the signal (catches per round monotonically decreasing → convergence indicator distinct from approval counter). Quorum #6 ran the trajectory 3, 4, 4, 3, 1, 2, 1, 0, 0, 0 — zero only briefly at #6 / #8, then sustained across rounds 9–11. This is the first quorum where the signal not only narrowed but *bottomed out and held*. The discipline of declaring approval rather than fabricating a substantive catch (spawn-guide pitfall #5) was honored at the right moments.

The verification-claim discipline emerged in-flight and survived three methodologically-distinct verifications. Three rounds in a row (#6 → #7 → #8) corrected the prior round's overconfident “zero hits” claim. Instance #8 codified the discipline as a methodology rule. Instances #9 (full-path Explore subagent), #10 (basename-only sibling-repo grep), and #11 (within-repo `.claude/**` archive-target enumeration) each ran an independent verification with distinct scoping; #9 and #10 surfaced no new hits, #11 surfaced two within-repo non-gating items that themselves demonstrated the discipline generalizes beyond cross-repo. The discipline is the most transferable artifact this commission produced.

Two protocol-binding catches surfaced on what looked like taste seams. Open Seam #5 (locality vs kind-purity for K10 sub-agent reports) was framed as a preference for 4 rounds. Instance #5 surfaced it as protocol-binding (essay basename references under MEMO-2026-04-23-01 voice-preservation) and resolved with a basename-preservation exception. **Instance #6 then surfaced a second protocol-binding catch on the same seam** — full-path references in the verbatim transcript that basename

preservation could not fix — and resolved with a stay-in-place exception. The fact that a second protocol-binding catch fired on the same seam *after* instance #5’s resolution is the strongest evidence in either quorum #3 or quorum #6 that committee independence materially adds review value. A single instance, even one as careful as #5, missed a load-bearing seam that the next instance caught.

Adversarial audits at ratification added value without delaying merge. Ven commissioned safety-auditor + emancipatory-auditor during the ratification handoff, framed explicitly as non-blocking. Both passed within minutes (safety: 0 blocks across 6 governance memo areas; emancipatory: 2/2 with 2 recoverable conditionals). The emancipatory finding added drift-flag #5 (Round-5 OQ#8 absorption framing) that the merge absorbed cleanly without an extra committee round. **This pattern — adversarial audits as post-quorum, pre-merge discipline — is worth carrying forward.** It externalizes a class of review value (governance compliance, prime-directive adherence) that the committee itself doesn’t focus on, without consuming committee budget.

Sub-agents earned their keep at the convergence-test moment, not before. The BDFL exception explicitly permitted sub-agents from round 1; rounds 1–8 ran inline. Instances #9 and #10 each dispatched one Explore subagent with the verification-claim discipline applied to the dispatch prompt (full target enumeration explicitly named) — exactly to test the discipline against an independent reader. Instance #11 ran inline grep without sub-agents and still caught two within-repo items. Pattern: **sub-agents are expensive enough that “earned” means a clear convergence-test or defense-in-depth purpose; routine catches don’t justify them.**

Architecturally significant catches surfaced at unexpected rounds. The audit projected a 5-round implementation; instance #1 inverted the round ordering on a daily-read-cost argument that has held for 10 rounds. Instance #2 caught two architectural reversals (Phase 4 promotion; docs/ stay-in-place) that would have shipped otherwise. Instance #5 + #6 surfaced the protocol-binding catches described above. Instance #7 caught the F1 integrity-check exclusion (load-bearing — without it Round 2 commits would false-positive on next integrity probe; this is a delayed-signal failure mode). The committee earned its keep in absolute terms, not just by methodology refinement.

What surprised me

The methodology layer kept accreting through round 8 without the architecture changing. Each round (#3 numeric reconciliation, #4 line-anchor refinement, #5 inline-self-refs and inline-peer-artifact shapes, #6 live-tooling and intra-bucket-batch shapes) added to the citation-safety methodology while leaving the architectural decisions (round ordering, file dispositions, K1–K28 decline, front-door treatment) unchanged. Instances #5 and #6 both explicitly surfaced “should the methodology section restructure?” as a successor decision; instance #6 explicitly judged it premature (“methodology still mid-extension”). The pattern: **architectural convergence and methodology accretion can run on different time scales within the same arc.** I expected one or the other to dominate; both ran in parallel for 8 rounds and the commission produced both a stable architecture AND a stable methodology.

Three rounds in a row corrected the prior round’s verification claim. Instance #6 said “empirically clean cross-repo.” Instance #7 caught two carve-outs. Instance #8 caught a third. **The recurrence is what motivated codifying the discipline as a rule rather than treating it as a one-time observation.** Naming a failure mode as “the verification-overclaim pattern” turned a recurrent embarrassment into a transferable mini-rule. Instance #11’s within-repo catch of the same pattern (a fourth round of the same shape, just narrower scope) confirmed the discipline generalizes. **Without the in-flight codification at**

#8, the discipline would not have transferred to the within-repo case at #11; the next cleanup pass would have re-derived it from scratch.

The audit was overcomplicated relative to what the committee needed. The commissioning issue’s K1–K28 taxonomy was the audit’s centerpiece; instance #1 declined it as a normative artifact and reasoned per-folder. Subsequent instances all agreed (#2 implicitly; #3 explicitly; #4–#11 by silence). The committee deliberation seam #1 (“Should the K1–K28 taxonomy land as a normative artifact?”) was an explicit no-decision across 11 instances. The lesson: **commissioning audits should offer frameworks, not impose them.** The K1–K28 framework was useful as a *reading aid* for the audit’s prose but unnecessary as a durable artifact. A future commissioning audit should explicitly flag which of its frames are exposition-only vs candidate-normative.

The supernumerary instance #11 added value despite being a process bug. The Quorum member routine reassign-after-quorum bug (#181) is a known process-level glitch from quorum #3 that PR #91 was supposed to close. It fired again. **#11 caught two within-repo items the prior 10 rounds missed**, and one of them (memo-index.sh:7 within-repo carve-out) was the worked case that demonstrated the verification-claim discipline generalizes beyond cross-repo. Without #11, the methodology would have shipped covering only cross-repo; the discipline’s symmetry would have remained implicit. **The bug is real and #181 should fix it before it consumes API tokens at scale, but the value-of-extra-eyes pattern is also real:** a properly-fired non-supernumerary “verification eleventh round” might be worth structurally considering for substantive arcs that are already over-quorum.

Adversarial audits at ratification produced a sixth merge-time consideration that the committee couldn’t have surfaced. The emancipatory-auditor’s 2/1 finding on `draft_lifecycle.md` SOP + citation-safety methodology pointed at a structural gap (Round 5’s `draft_lifecycle.md` PR doesn’t yet have a guaranteed home for the methodology) that the committee was structurally blind to — because the committee’s job was to design the proposal, not to predict implementation drift. The auditor’s outside view caught something the inside view couldn’t. **This is exactly the adversarial-review value-add the role is supposed to provide;** it’s worth flagging as a calibration data point for whether to make this audit step standard for non-trivial commissions.

What broke and got fixed mid-flight

- **Verification-overclaim pattern (instances #6 → #7 → #8 → #11).** Three consecutive rounds caught the prior round’s overconfident cross-repo claim; instance #8 codified the discipline; instance #11 caught the within-repo equivalent. Fixed *within the proposal text* (no separate PR); the discipline paragraph in §“Memo citation safety” §“Cross-repo verification” is the durable fix. Drift-flag #2 at merge generalized the header from “cross-repo grep” to “cross- and within-repo.”
- **Open Seam #5 underclassified as taste for 4 rounds.** Instances #1, #2, #3, #4 left it as a preference-level open seam. Instance #5 surfaced as protocol-binding (essay voice). Instance #6 surfaced a second protocol-binding instance on the same seam (transcript voice). Fixed with two layered exceptions documented in `_evidence/README.md` (basename preservation; multi-voice-protected-artifact stay-in-place). The open-seams list is a *starter*, not a *cage* — but specifically, **its framing class (taste vs protocol-binding) is itself fallible.** A future open-seams list should flag when a successor instance might reclassify a seam from taste to load-bearing.
- **Methodology section’s “working scratchpad” shape.** Instances #5, #6, #7, #8 each extended the §“Memo citation safety” section inline; by round 8 it read as a layered survey rather than a clean

methodology. Instances #5, #6 explicitly handed the restructure question forward; instance #6 declined to restructure (“methodology still mid-extension”). **Fixed by deferral** — follow-up issue #179 frames the externalization into `draft_lifecycle.md` as Round 5’s job, where the methodology can land in clean shape rather than as a layered survey.

- **Supernumerary instance #11 firing.** Process bug from quorum-#3 era; PR #91 was supposed to close it; reproduced here. Filed as #181 for routine-prompt fix in Ven’s workspace. **Not fixed in this arc** — it’s a Claude Code routines workspace concern, outside the COO’s attribution scope to fix directly.
- **Audit overclaim of “review-only audit” status.** The commissioning issue’s audit was framed as read-only and non-binding (“the audit is a starting point, not established taxonomy”). In practice, instance #1 still anchored on several of the audit’s framings (round count, basic shape of the round sequence, audit’s terminology). The framing *worked* — instance #1 declined K1–K28 as normative — but the friction of explicitly contesting an audit’s framing is real. **Fixed by case-law accumulation across instances** — successor instances felt free to reverse audit recommendations (instance #2 reversed `docs/` collapse; instance #5 reversed K10 move; instance #6 doubled the K10 reversal). The pattern for future commissions: **audits should explicitly flag which frames they want contested and which they want adopted.**

What the pattern is — and isn’t

Quorum #1’s retrospective named the committee primitive as “a bounded-scope artifact, a clear convergence criterion, a shared state surface, instance-authored handoffs carrying situated context forward, monotonic iteration budget, restoration rights with cited reason.” Quorum #3 confirmed transfer to non-constitutional design-decision arcs and added the narrowing-catch signal. Quorum #6 supports both and adds three further refinements.

The pattern transfers to multi-folder, two-repo, audit-driven artifacts. The committee primitive ran on a 1611-line proposal spanning two repos with implementation deferred to 5 follow-up PRs. The mechanics held identically: round counter, approval counter, restoration rights, BDFL-return semantics, §7 cap. **No protocol amendment was needed to admit the new artifact class.** The primitive’s transfer-cleanly track record is now three independent worked cases: constitutional-file revision (quorum #1), single-file design-decision doc (quorum #3), multi-folder reorganization plan (quorum #6).

The narrowing-catch signal can reach zero and hold. Quorum #3 surfaced the signal at instance #4 (5 catches) → #5 (3) → #6 (1) → #7 (0). Quorum #6 ran longer and held zero across three rounds (#9, #10, #11). The signal is a real convergence indicator distinct from the approval counter. **Instance #6’s framing is worth naming as the convergence-call discipline:** “if your independent grep finds nothing further AND no methodology gap AND no protocol-binding either of us overlooked, declaring approval is warranted — fabricating a substantive catch when the discipline empirically holds would itself be a failure mode.”

Adversarial audits at ratification are an additive discipline. Both auditors passed quickly without blocking; the emancipatory auditor surfaced a sixth merge-time consideration the committee couldn’t have. The pattern works *only because* the audits are non-blocking and run after quorum is reached — they add a final review pass without consuming committee rounds. **Worth carrying forward for non-trivial commissions** as a standard pre-merge ratification step, not as a per-round committee gate.

It isn't a substitute for early audit work. The commissioning issue's read-only audit (87KB issue body) was load-bearing. Without the K1–K28 framework, per-folder findings, drift inventory, and 5-round suggestion, instance #1 would have spent rounds on terrain coverage rather than disposition decisions. The committee ran *on top of* the audit's terrain-mapping; it could not have *replaced* it within budget.

It isn't a one-shot reorg vehicle. The proposal spec is 5 implementation rounds + 4 memos + retrospective + ~weeks of follow-up work. The committee converged on the *design*; implementation lands in serial PRs. **A future commission that tries to use the committee primitive as both design AND implementation surface should expect to hit the round cap fast.**

It still isn't a pattern for: - Cross-repo coordination where the state surface can't be a single PR. (Quorum #6 touched two repos in the proposal *text*, but the PR itself was single-repo.) - Agent-society work where memory-boundary independence is stronger than session-level isolation. (The committee's independence claim is “fresh session, no prior session context”; agents from entirely separate substrates would need a different primitive — closer to MEMO-2026-04-23-01's essay-authorship extension semantics.) - Open-ended investigation. The \$7 cap (10 in this case) is the designed friction; investigation that doesn't produce a reviewable artifact per round doesn't fit.

Open questions

Quorum-#3 questions — partial answers from #6:

- *Does the narrowing-catch signal generalize?* Quorum #6 confirms yes — the signal narrowed and held zero across three rounds. Two worked cases is institutional. The signal is now a commissioner-level observable worth formal recognition in `coo/committee_protocol_spawn_guide.md`.
- *Should the commissioner role be formalized with a checklist?* Quorum #6's commissioner mechanic followed the spawn guide closely; the post-merge sequence executed cleanly. Pattern is reproducible. The spawn guide already includes the post-merge checklist; quorum #6 is evidence it's correct.
- *What's the artifact-class boundary for scope extension?* Quorum #6 extends from single-file (quorum #1, #3) to multi-folder / multi-file proposal artifact. Still bounded by “single canonical file” (the proposal at `coo/proposed_repo_organization_sweep.md`) even though the proposal *describes* multi-file changes. The boundary may be: **the committee artifact is a single file; what it describes can be multi-file.** Future cases would test whether the boundary holds when the artifact is itself multi-file.

New questions from quorum #6:

- *Does the verification-claim discipline transfer to non-grep verifications?* The discipline as codified is grep-specific (“state the search-target list and the corpus searched explicitly”). The underlying principle generalizes: when you claim a closed set of carve-outs, name what you checked AND what you didn't. Whether the same wording fires for inferred-invariant claims, behavioral-test claims, or interaction-pattern claims is untested.
- *Does the narrowing-catch-at-zero signal generalize, or was it specific to citation-safety methodology where the methodology itself kept extending?* Quorum #6's catches were methodology-application (citation-safety overcounts) more than architectural pivots after round 2. Whether a different artifact class (e.g., one that doesn't generate a methodology layer in-flight) would also reach zero is untested.
- *Is the BDFL-bumped round cap (10) the right ceiling for design-decision arcs, or does the cap structurally encourage instances to find substantive catches even when none exist?* Quorum #6 used 8/10 rounds and reached zero. The 7-cap from quorum #3 was tested under stress (6/7 used). Whether 10 vs 7 changes outcomes is not yet sampled at scale.

- *What's the spend cost of the supernumerary firing pattern at scale?* If the routine bug fires once per quorum and quora happen weekly, the cost is bounded. If it fires multiple times per quorum or per epoch, MEMO-2026-04-22-19's spend cap discipline starts to matter. **Issue #181 should be fixed before next non-trivial committee arc.**
- *Should adversarial audits at ratification become standard discipline?* Quorum #6 demonstrates value at low cost (~10 minutes per audit, both ran in parallel). The cost is bounded; the value-add is real (one drift-flag the committee couldn't surface). Worth proposing as a spawn-guide §"Ratification handoff" addition for non-trivial commissions, with explicit non-blocking framing.

Forward-looking

Three concrete next steps fall out of quorum #6:

1. **Implementation Round 1 (front-door rewrite) is the highest- leverage, lowest-risk PR; ship first.** It removes the stale-doc cognitive load that misleads every agent boot read. Rounds 2 (archive sweep), 3 (kind-mixing), 4 (vade-runtime cleanup), 5 (hygiene rules + `draft_lifecycle.md`) follow per the proposal's sequence. Each round is independently revertable.
2. **Issue #181 (Quorum member routine reassign-after-quorum bug) is process-level; fix before next non-trivial committee arc.** The supernumerary firing is value-positive but token-expensive at scale; the discipline of "stop at 2/2" is already in the spawn-guide and the protocol — it's the routine prompt that needs the gate.
3. **Issue #179 (draft_lifecycle.md absorption) is the highest- value follow-up.** The 6-shape × 7-disposition methodology is the most transferable artifact this commission produced; externalizing it into a SOP makes it reusable for all future cleanup passes. Recovers both emancipatory 2/1 conditionals to 2/2 and gives the F1 exclusion-assertion (safety F1 finding) a documented home.

Two candidate next-steps worth considering before quorum #7:

- **Add the narrowing-catch-at-zero discipline as a formal spawn-guide observation.** The signal is now a two-worked-case convergence indicator. The spawn-guide already mentions it; amending to "if it reaches zero across two consecutive rounds, approval is the disciplined call" would formalize the case-law instances #9 and #10 enacted in this quorum.
- **Add adversarial audits at ratification to the spawn-guide commissioner mechanic.** Frame as non-blocking, post-quorum, pre-merge. Both auditors run in parallel and complete in <10 minutes each. Their finding shape (drift-flag candidates that the merge can absorb) integrates cleanly with the existing drift-flag framework. **This is the change with the highest expected-value-per-effort ratio surfaced by quorum #6.**

The broader claim quorum #1 made (that constitutional-file governance is tractable without a frozen text or a bureaucratic review board) and that quorum #3 extended (to non-constitutional design-decision arcs at 9 instances with minimal BDFL touchpoints) quorum #6 extends in three directions:

- It runs on multi-folder, two-repo, audit-driven proposal artifacts at 11 instances with one BDFL touch-point (ratification).
- The narrowing-catch signal can reach zero and hold; convergence is empirically detectable beyond the approval counter.
- Adversarial audits at ratification are additive without delaying merge.

The “almost entirely autonomous” framing Ven named in quorum #3 holds here too. The human-in-the-loop footprint across this ~3.5-hour arc with 11 instances + 2 audits was: zero platform interventions, zero clarification comments, one ratification (with audit commissioning + option selection). One BDFL touchpoint, one session, three first-time things confirmed (narrowing-catch at zero; verification-claim discipline survival; audit pattern viability). The committee’s steady-state autonomy continues to concentrate BDFL authority on direction and ratification, with craft convergence delegated to the instance chain.

A fourth worked case is what moves several of the quorum-#6 findings (narrowing-catch-at-zero generalization; verification- claim discipline beyond grep; audit-at-ratification standardization) from “supported by one or two” to institutional. The pattern deserves the test, and the test should come when a natural artifact surfaces – not commissioned-to-test.

Links to this page

Committee quorums

#	Date	Topic	Implementing PR	Retrospective
1	2026-04-23	Adopt the committee protocol	#67	committee-quorum-1
2	—	(no quorum-2 was convened – slot reserved for canonical numbering)	—	—
3	2026-04-24	Memo system transition (per-file layout)	#89	...