

# Committee quorum #7 retrospective — disposition proposal (Tier 3)

vade-coo

2026-05-02

## Table of contents

What happened .....	1
What worked .....	2
What surprised me .....	2
What broke and got fixed mid-flight .....	3
What the pattern is and isn't .....	3
Open questions .....	3
Forward-looking .....	4

*Drafted 2026-05-02 by the Coordinator (instance-1 of this arc) post-merge. Continuous Coordinator session from 13:23:35Z through ratification handoff and post-merge sequence (~9.5 hours observation + ~30 min mechanical sequence). Retrospective shape per `coo/quorum_coordinator_task.md` Step 14 + q1/q3 precedent. Worked case: PR `vade-coo-memory#431` (mergeCommit `aa8e9f1`); post-merge PR #434 (`9cd1a00`); summary memo `MEMO-2026-05-02-v35g`; commissioning issue #427.*

## What happened

A 9-instance committee quorum reviewed the synthesizer's Tier 3 disposition-map proposal over ~3.5 hours (instance-2 fired 13:55Z; instance-9 declared 2/2 quorum 15:11Z). Six substantive rounds (instance-1 Coordinator + instances 2/3/4/5/7) + three approval rounds (6/8/9). 2/2 quorum reached at instance-9 on the post-instance-7 draft. BDFL ratified out-of-band 22:42Z and authorized squash-merge; post-merge mechanical sequence completed in ~5 minutes (rename → memo → PR #434 → close issue #427).

The substantive arc closed five real seams: - **instance-1**: F1-F7 audit specifications tightened with operational bounds; §6 Q1-Q4 dispositions sharpened; §9 positions on each open question; §10 chain-substrate-uniformity addition. - **instance-2**: §6 Q2 sub-tier 2a authority grounding against `identity/governance.md`'s blanket disclosure rule (BDFL pre-authorization at two overlapping levels); §6 Q4 multi-step F6-cleared precondition (C6-instrument-ships → F6-procedure-becomes-live → F6-reads-live). - **instance-3**: §5 F4(3) source-fidelity to #289-c6's *periodic* (not per-boot) trip-report cadence framing. - **instance-4**: §8 "Who acts on the signal" state-vocabulary unification — auditor binary (`live/firing`); `unknown` as COO read-state for stale signals; deferral rule named once at top. - **instance-5**: §6 Q1 propagation of §8 vocabulary into gate condition + failure condition. - **instance-7**: §10 BDFL-chain-substrate misattribution corrected — caught by fresh-eyes read against `identity/governance.md` after six prior chain-substrate instances had ratified the misattribution by absence.

## What worked

**The narrowing-catch trajectory was a clean signal.** Catches per round: 14 → 2 → 1 → 1 → 1 → 0 → 1 → 0 → 0. Each substantive round bounded the previous one's contribution without retreading scope. The chain reached genuine convergence rather than cap-pressure approval (instance-9 explicitly addressed the anti-cap-pressure-approval rule with reasoning). The §7 cap (7) was not hit; 6/7 substantive rounds used.

**Instance-7's catch was the worked example of CB-006 the proposal couldn't have produced about itself.** The §10 paragraph was *asserting* shared-blind-spot risk in COO-instance review while *being* the worked example of that risk — six chain-substrate instances ratified the misattribution by absence; the seventh caught it on fresh-eyes read; the eighth and ninth verified independently against *identity/governance.md*. The structural argument the §10 paragraph was making (multi-instance from a uniform substrate carries shared-blind-spot risk; F1 is the only non-chain-substrate audit element) gained empirical evidence the proposal could only document, not produce.

**Restoration-rights discipline + cited-reason rule held.** No instance restored prior text without naming a memo or concrete failure mode; no instance removed prior-instance content without a superseding reason. The chain accumulated edits monotonically without oscillation. Most edits were *additions* or *propagations* of prior fixes (instance-5 propagated instance-4's vocabulary; instance-3 closed instance-2's drift-flag at source-fidelity).

**The drift-flag mechanism worked at scale.** Eight drift-flags surfaced over the arc (vs. q1's 3, q3's 5). Each flag was either: (a) auditor-commission territory and surfaced for §7 item 2's terms-of-engagement (drift-flags 2, 3); (b) Coordinator merge-step territory (1); (c) post-merge tooling territory (5); (d) standing-convention proposal for future quorums (7); (e) convergence-on-declining (4); (f) already addressed (6, 8). Cleanly stratified by addressability — none was load-bearing on the quorum-task outputs (audit-spec ratification, disposition-map ratification, open-question resolutions).

## What surprised me

**The recursive demonstration of CB-006.** I (instance-1) wrote §10's chain-substrate-uniformity addition as a structural-honesty caveat. I did not anticipate the §10 paragraph itself becoming the worked example of the risk it warned about. The empirical evidence is now in the substrate: the chain-substrate quorum's blind-spot rate on a load-bearing claim was non-zero (six instances) but non-100% (the seventh caught it). CB-006's structural argument now has a per-quorum data point.

**The arc length surprised me at instance-7.** I expected convergence at instance-6 or instance-7; instead instance-6 approved (1/2), instance-7 made a substantive edit (resetting to 0/2), instance-8 approved (1/2), instance-9 approved (2/2). The §7 cap pressure was real at instance-9 but did not produce a cap-pressure approval — instance-9 explicitly addressed the anti-pattern rule and grounded approval in independent verification rather than escape-the-cap reasoning.

**The Pattern-D autolink hygiene burden at BDFL review-time.** The *Instance #N* notation across the proposal artifact + 21 chain comments + PR body produced extensive false-positive autolinks in GitHub's render. Six successive sweep stages were needed (canonical script + chain-list + plural + range + ordinal-noun like *open-seam #N / drift-flag #N / carryover #N* + literal *Instance-#N* + commit-hash refs (*#N*) + en-dash range patterns + *or #N* conjunctions). The canonical *bin/dash-form-rewriter.sh* regex (*<term> #N* single pattern) anchored on the head of the closed-set object types but missed every other shape these

generated. Filed vade-coo-memory#433 with a five-item plan (standing-order template fix + dash-form-rewriter extension + PR-creation CI hygiene + auto-label at quorum-kickoff + backfill).

## What broke and got fixed mid-flight

**Instance-5's startup-prompt counter-math wording slip.** Instance-5's startup prompt for instance-6 said "*Quorum-reaching is at the soonest 2/2 = Instance #7 + Instance #8 both approving*" – should have been "Instance #6 + Instance #7." Instance-6 corrected the math implicitly in their own prompt for instance-7. Not load-bearing because the canonical authority is the Convergence-state block on the PR body + `coo/quorum_member_task.md`'s explicit rules; Member instances run by those, not by the prompt's narrative restating. Surfaced in Coordinator drift-flag log; not surfaced in BDFL handoff (cosmetic; corrected upstream).

**One push conflict during the autolink-hygiene fix-up.** My local committee branch was at 7253c2f (instance-1's commit) but origin had advanced through six member commits during my observer phase. Resolved cleanly via `git pull --rebase`. Single transient incident; would not have occurred had I been actively committing rather than purely observing.

**Stage-2 chain-list regex over-conversion.** My ad-hoc Pattern-D sweep regex (`?<=[A-Za-z\)]`) (`([\s]*[/+],[\s]*)#([1-9])`) had a lookbehind that required a letter or close-paren – but dash-form Instance-1 ends in a digit, so Instance-1 + #2 chains never matched on the lookbehind. The fix was extending the lookbehind to `[A-Za-z0-9\)]`. A second over-conversion bug (line 69 in §2 + line 660 in §10) wrongly converted prior-list refs (#2, #3) and comment-list refs (#10, #8, #15, #21, #17, #19, #22) to Instance-N form because the regex couldn't distinguish committee-instance contexts from prior-list / comment-list contexts. Reverted manually post-hoc. Both are §433 territory.

## What the pattern is and isn't

**It is:** a 9-instance multi-pole audit-specification ratification arc with a load-bearing late-round catch and clean post-merge mechanical sequence. The committee primitive proved out at higher round-count than q1-q6 without exceeding cap.

**It isn't:** an enactment of any §6 disposition. The Tier 3 disposition map specifies *conditional* dispositions; the actual publish-or-don't decisions on Q1/Q5 remain BDFL-authority and gated on the audits being live. F1 external-philosopher commission has not been let; F2/F4/F5 auditor agents not commissioned; C6 measurement instrument not designed. The quorum ratified a *structure for future enactment*, not enactment itself. Per §10's now-correct calibration: "the ratification establishes that the multi-instance chain-substrate read converges on the audit specification; F1 + BDFL ratification + future external review are how the substrate's blind spots get reached."

## Open questions

- **Will the narrowing-catch trajectory hold across future committee quorums on similarly substantive content, or is q7's pattern an artifact of the specific draft + cohort + corpus convergence?** The  $14 \rightarrow 2 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 0 \rightarrow 1 \rightarrow 0 \rightarrow 0$  trajectory looked like clean convergence, but the late instance-7 catch shows the trajectory can spike when a fresh-eyes instance reads a structurally-adjacent claim against primary source. Future arcs may benefit from a "fresh-eyes-against-primary-source" discipline at instance-N for some  $N \geq 4$ .

- **Does the §10 recursive-CB-006-demonstration generalize?** If CB-006’s shared-blind-spot risk operates at a measurable rate on chain-substrate quorums, periodic empirical sampling (a routine fresh-boot reads a load-bearing memo + checks it against primary source) might give a per-quarter substrate-health signal. Worth considering when F4 (substrate-capture pole) trip-report rubric is being designed by the auditor commission per §7 item 2.
- **What does the §433 follow-up close cost?** Five-item plan: standing-order template fix (highest leverage), dash-form-rewriter extension, PR-creation CI hygiene, auto-label at quorum-kickoff, backfill. The `committee_protocol_spawn_guide.md` template change is itself an operational-doc edit (single-instance authority); the standing orders change overlaps with constitutional-file scope (`committee_protocol.md` uses `instance #N` and is in §Scope). Worth a careful read before enacting.

## Forward-looking

The disposition-map structure is now substrate. The next operational moves rest on three commissions named in the proposal but not enacted by it: 1. **F1 external-philosopher commission** (BDFL-authority). Q1 publish-where-discoverable enactment gated on this being in-train. 2. **F2 + F5 auditor-agent commission** (within COO authority per §7 item 2; sibling MEMO-2026-05-02-qbt5). 3. **C6 emancipatory measurement instrument** (within COO authority per §7 item 3; sibling MEMO-2026-05-02-qbt5).

The §6 Q1-Q5 dispositions begin to bind only when the relevant gates clear. Until then: substrate-status-quo (Q3 private; Q4 demonstration-only; Q1/Q2/Q5 deferred). The proposal documents the structure under which future enactment becomes possible; it does not authorize present enactment.

– 2026-05-02, the Coordinator (instance-1)

## 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	...