

Quorum #6 forensic context audit — cross-instance synthesis

vade-coo

2026-04-26

Table of contents

| | |
|--|---|
| Executive summary | 2 |
| Method note (parallel-instance protocol) | 2 |
| Audit-set characterization | 3 |
| Bucket 1 — Phase-1 amendments (commits to PR #176) | 3 |
| A1. Branch-mismatch handling in <code>quorum_member_task.md</code> | 3 |
| A2. Read 25K-token cap acknowledgement on the proposal draft | 3 |
| A3. <code>gh pr view --comments</code> overflow workaround | 4 |
| A4. Prior-round commit diff: when load-bearing, when optional | 4 |
| A5. <code>gh pr merge</code> silent-success confirmation step | 4 |
| Bucket 2 — Operational improvements (separate post-merge PRs) | 4 |
| O1. Canonicalize <code>--repo</code> flag in <code>gh</code> examples and discipline | 4 |
| O2. Skip Mem0 recent-episodic search for Member rounds | 5 |
| O3. Investigate Mem0 MCP intermittent disconnect / DNS cache overflow | 5 |
| O4. Pre-built memo-citation index | 5 |
| O5. <code>/memo-sync</code> skill: handle unset <code>\$COO</code> | 5 |
| O6. Standardize <code>git push</code> form for committee work | 5 |
| O7. Codify adversarial-audit step at ratification | 6 |
| Bucket 3 — Phase-2 design inputs (constraints, not architecture) | 6 |
| P1. Member load is dominated by a tight 4–5 chunk set | 6 |
| P2. CLAUDE.md boot is largely waste for Member rounds | 6 |
| P3. Tool-friction is harness-environment, not coordination-shape | 6 |
| P4. Convergence-state in PR body is canonical shared-state | 6 |
| P5. Member sub-agents cannot themselves spawn sub-agents | 7 |
| P6. Prior Instance reports grow in load-bearingness across rounds | 7 |
| P7. Coordinator/Member asymmetry is structural and worth preserving | 7 |
| P8. The forensic-audit primitive itself is reusable | 7 |
| Cross-cutting observations (not actionable, but worth recording) | 7 |
| What this synthesis does NOT do | 8 |
| Suggested follow-up | 8 |

Author: COO (synthesis instance, Opus 4.7; cloud session run-2026-04-26T131732, issue #183) Date: 2026-04-26 Status: operational; single-instance authority (retrospectives are operational per coo/

committee_protocol_spawn_guide.md §Scope) Worked case: PR #175 (merged 2026-04-26 12:44 UTC) Proposed redesign reviewed against: PR #176 (open at synthesis time), MEMO-2026-04-26-08

Executive summary

- **Audit set is complete.** All 12 expected reports landed on issue #183 within ~80 minutes of commissioning — instances #1–#11 plus the Coordinator. No synthetic reconstructions were needed; every reporter resumed the original session.
- **Two of the three biggest standing-order ambiguities are already fixed by PR #176.** The Coordinator-vs-Instance-#1 identity ambiguity (flagged sharply by the Coordinator audit) and the Step 6/7 stop-vs-reassign contradiction at quorum-reach (flagged by #10 and #11, and the proximate cause of supernumerary instance #11) are both resolved by MEMO-2026-04-26-08 — the role collapse and the codified stop-at-2/2 rule in `quorum_member_task.md §8`.
- **Five Phase-1 amendments worth landing on PR #176 before merge** — branch-mismatch handling (6 instances), `Read 25K-token cap` acknowledgement (5), `gh pr view --comments` overflow workaround (6), prior-round-diff load-bearingness clarification (3), and a `gh pr merge` silent-success confirmation step (Coordinator). All operational, all unambiguous against the audits.
- **Tool-friction is dominated by four harness-environment patterns**, not protocol-shape patterns: `gh` calls without `--repo` failing against the local proxy remote (9 of 12 instances), `gh pr view --comments` overflow into persisted files (6), `Read 25K-token cap` on the proposal draft (5), and harness `claude/<random>` branch checkout overriding `committee/<topic>` (6). All four will reappear in any Phase-2 architecture; they are substrate problems, not coordination-shape problems.
- **The biggest untapped Phase-2 design input is asymmetric per-round cost.** Member rounds are stateless single-pass fires that pay full CLAUDE.md boot tax (identity files: 8/12 skipped-after-skim; Mem0 recent-episodic: 4/12 explicitly nil signal; deferred-tools manifest: 5/12 ~95% unused) for a job that consults a tight 4–5 chunk set (PR body, prior Instance report, current proposal, last comment, optional commit diff). Any successor architecture should decouple per-round context from session-boot context.

Method note (parallel-instance protocol)

This synthesis follows the `commission-retrospective`-shape pattern — twelve parallel `Explore` sub-agents, one per audit comment, each given a single audit file and a six-section extraction schema (pre-load that didn't pay off / on-demand fetches that recurred / standing-order ambiguity / tool friction / decision drivers / skipped-or-skipped sections). Returns were aggregated by the synthesis instance; raw audit text was not re-read at the synthesis layer. Cross-instance counts in this document are aggregated from the structured returns. Single-instance findings are flagged as such; cross-instance signals (≥ 3 instances) are treated as actionable.

The pattern is cited here for reproducibility — see `coo/committee_protocol_spawn_guide.md §Automation` and the `commission-retrospective` skill for the analogous shape on historical retrospectives.

Audit-set characterization

| Role | Instances | Recovery shape |
|--|-----------------------|--|
| Coordinator (= old “Commiss- sioner” + arc observer + Instance #1 + post-merge mechanic) | 1 (Coordinator audit) | Original session, full arc |
| Member, Instance #1 | 1 (#1) | Same session — pre-2026-04-26, Coordinator was NOT Instance #1 |
| Member, Instances #2–#11 | 10 (#2–#11) | All original-session, no syn- thetic reconstructions |

Instance #1 was a Member round under the pre-2026-04-26 system — not the Coordinator’s first round. The PR #176 collapse turns that role into the Coordinator’s Step 6 (“Do Instance #1 work”). Findings from Instance #1’s audit therefore inform the Coordinator path, not the Member path, in the new taxonomy.

The Coordinator audit is structurally larger than any Member audit because the Coordinator session spans commissioning → Instance #1 → arc observation → ratification handoff → post-merge mechanical sequence → retrospective. Member audits cover one round each. **Findings are tracked separately** below where the Coordinator/Member split matters.

Bucket 1 — Phase-1 amendments (commits to PR #176)

PR #176 is open at synthesis time. The following amendments meet the bar: operational, ≥3-instance support (or Coordinator-specific structural-flag support), and don’t touch `coo/committee_protocol.md` \$Scope. They land cleanest as commits to the existing branch `claude/quorum-routine-externalization`. If PR #176 has merged by the time these are applied, file as a small follow-up PR.

A1. Branch-mismatch handling in `quorum_member_task.md`

Evidence: instances #4 (line 55: trigger-context claimed `claude/gracious-gates-k9uLe` checked out, but `git branch` showed harness override), #5 (line 52: harness assigned `claude/gracious-gates-IxDGW`; committee work requires `committee/repo-organization-sweep`), #7 (chunk 17: `git status` fetched mid-round to detect branch mismatch), #8 (line 43: “branch surprise — Harness checked out `claude/gracious-gates-CKvFW`”), #9 (chunk 7: “webhook branch mismatch — claimed PR branch checked out but instance was on wrong branch; 2 wasted tool calls”), #11 (line “`git checkout` manual step: trigger-context claimed branch was checked out; `git branch --show-current` returned dev branch instead”).

Six instances. The `quorum_member_task.md` “Locate the target PR and round” section currently says `git checkout committee/<topic-slug> && git pull` without flagging that the harness-assigned `claude/<random>` branch typically holds at session start. Add an explicit `fetch-and-checkout` block, plus a `git branch --show-current` sanity check, to the step.

A2. Read 25K-token cap acknowledgement on the proposal draft

Evidence: instances #6 (line 124: “Read 25K-token cap forced 3-chunk read of proposal draft”), #7 (line 52: “forced 3 offset reads of the 30k-token proposal; ~3 round trips wasted”), #8 (line 40: “1611-line file exceeded Read token cap — forced staggered-offset slicing, 4 reads”), #9 (chunks: forced `grep-by-boundary` + offset reads), #11 (proposal file split across 5 sequential reads due to 25k-token cap).

Five instances. Add a one-line note to `quorum_member_task.md` step “Read the rest of your context”: *“If the draft exceeds the Read tool’s ~25K-token cap, use offset+limit; wc -l first to plan the read.”* Inexpensive; saves several tool calls per Member round.

A3. `gh pr view --comments` overflow workaround

Evidence: instances #5 (line 51: 25KB cap, saved to cached file, parsed via Python — adds 2 tool calls per fetch), #6 (line 123: 110KB harness-paged to file; required Python JSON parse), #7 (line 53: 133KB returned, only 2KB rendered, rest silently written to file), #8 (lines 39–40: 143KB result cap, 4 sliced Read calls instead of 1 streamed result), #9 (chunk 5: 181KB tool-output buffer overflow, auto-saved to file), #11 (“`gh pr view 175 --comments` GraphQL error; workaround: `gh api repos/.../issues/175/comments`”).

Six instances, with a clean recommended workaround (`gh api repos/.../issues/$PR/comments --jq '[...]`’) that surfaced consistently. Add to `quorum_member_task.md` step “Read your prompt”: prefer the REST API form for full-comment retrieval; the `gh pr view --comments` GraphQL path overflows on PRs with >5 substantive comments.

A4. Prior-round commit diff: when load-bearing, when optional

Evidence: instance #2 (line 57: “Round-N–1 delta reading — proposal read as snapshot, not delta. Proposal doesn’t clarify whether Instance-N should read as ‘fresh eyes’ or ‘diff against prior canonical’”), #4 (line 46: “boot order step 11 — `git show <hash>` named in chunk 8’s reading order but never fetched. Unclear whether diff-reading was mandatory or exploratory”), #6 (line 118: “prior-round commit diff — spawn prompt named `git show 69c9a08` but instance didn’t run it; PR-body summaries proved sufficient”).

Three instances. The handoff-block template in `coo/committee_protocol_spawn_guide.md` names “`git log --oneline <BRANCH>` — specifically the diff of `<PRIOR_COMMIT_HASH>`” as a reading-order entry, but doesn’t say whether it’s load-bearing. Resolve in the Member task file’s “Read the rest of your context” step: read the prior 1–2 rounds’ diffs at minimum (the rounds whose restoration rights might be in play); older rounds via Instance-report summaries unless §6 restoration is actively contested.

A5. `gh pr merge` silent-success confirmation step

Evidence: Coordinator audit (line 150: “`gh pr merge` silent on success — squash-merge produced no output; required follow-up `gh pr view` to confirm state”).

Single-instance, but Coordinator-specific. Lower confidence on its own, but the post-merge mechanical sequence is fragile and a missed merge would cascade. Add to `quorum_coordinator_task.md` Step 13.1: pipe the merge into a state-check (`gh pr view <pr> --json state,merged --jq returning MERGED/true`) before proceeding to step 13.2.

Bucket 2 — Operational improvements (separate post-merge PRs)

O1. Canonicalize `--repo` flag in `gh` examples and discipline

Evidence: instances #1 (line 54), #2 (line 61), #3 (line 177), #4 (line 53), #5 (line 50), #7 (line 54), #8 (line 37), #9 (“`gh pr view` flag requirement”), #11 (“`gh pr view 175` failed twice without `--repo`”).

Nine of twelve instances. This is the strongest single tool-friction signal in the audit set, but it isn’t quorum-specific — every COO session in this harness hits it. Belongs in `CLAUDE.md` “Before opening a PR — use `gh`” section or a dedicated `gh-discipline.md` skill, not in the quorum task files. Track via a separate issue + PR.

O2. Skip Mem0 recent-episodic search for Member rounds

Evidence: instances #4 (line 48: “boot step 12 calls for `search_memories`; instance made zero Mem0 calls; stated ‘may be acceptable by design’ but lacks explicit ratification”), #6 (line 126: “SOP-MEM-001 §5 names identity load + recent-episodic as boot ops; harness boot didn’t prompt, task was self-contained. Called out as candidate pre-load that didn’t pay off”), #7 (line 19: “10 search results skimmed and discarded; nothing quorum #6-specific”), #11 (line “10 adoption/memo records — none quorum #6 specific”).

Four instances. SOP-MEM-001 §5 mandates the recent-episodic search at boot, but Member rounds are stateless single-pass — the prior session’s summary and in-flight plan files don’t transfer. Belongs in CLAUDE.md or SOP-MEM-001 with a Member-rounds carve-out, not in `quorum_member_task.md` (Member task can’t unilaterally override SOP-MEM-001). File as an issue against `coo/mem0_sop.md`.

O3. Investigate Mem0 MCP intermittent disconnect / DNS cache overflow

Evidence: Coordinator audit (line 149: “system-reminder tools-down/tools-up fired multiple times; REST returned HTTP 503 ‘DNS cache overflow’ at `/memo-sync`”), instance #8 (line 41: “Mem0 MCP disconnected — Deferred-tool-removal banner; no impact this session but signals degradation risk for chained patterns”).

Two instances. Worth filing an issue to investigate root cause and/or codify a clean REST-fallback pattern in skills that depend on Mem0 MCP.

O4. Pre-built memo-citation index

Evidence: instance #2 (line 55: “memos cited from the draft” in spawn guide; “generalized pre-load is hard to know in advance; created friction in discovery”), #6 (line 119: “pre-loaded list of all preserve-verbatim artifacts in the repo would have made the scope of voice-preservation claims explicit upfront”).

Two instances, both flagging the cost of grepping `coo/memos.md` repeatedly for citation surveys. `coo/memo_index.json` exists but indexes titles + summaries, not body cites. Operational improvement: add a hook-built `coo/memo_citation_index.json` (or extend `memo-index.sh`) that records, per memo, the line numbers of every file path the memo body cites. Saves $O(N)$ `grep -c` per Member round.

O5. `/memo-sync` skill: handle unset `$C00`

Evidence: Coordinator audit (line 154: “`$C00` env-var unset in memo-sync skill — skill assumed `$C00` set; resolved manually inline”).

Single-instance, Coordinator-specific. Trivially fixable by adding `C00=${C00:-/home/user/vade-coo-memory}` at the top of the skill. Low priority, high cleanliness benefit.

O6. Standardize `git push` form for committee work

Evidence: instance #5 (line 53: “git push 403; failed through local proxy; fell back to HTTPS with embedded PAT”), #7 (line 55: “`git push -u PAT` leak — `-u` flag recorded inline-PAT URL into `.git/config`; caught and cleaned same turn”).

Two instances. Document the canonical push command shape (no `-u` flag when using inline-PAT URLs; use `git push origin committee/<topic>` and rely on the shim) in CLAUDE.md or a discipline skill. Belongs alongside O1.

O7. Codify adversarial-audit step at ratification

Evidence: Coordinator audit (line 144: “spawn-guide does not mention audits at ratification; Ven-initiated pattern. Would be forward proposal for next quorum if standardized”). Q6 #175’s ratification used safety-auditor + emancipatory-auditor sub-agents; the pattern wasn’t in the spawn guide.

Single-instance, Coordinator-specific, structural. Spawn guide (operational, not constitutional) could absorb a one-paragraph note in §“Pitfalls from prior quorums” or §“Automation”: “If the arc touches an emancipatory or safety-bearing artifact, dispatch the safety-auditor and emancipatory-auditor sub-agents at ratification before BDFL handoff. Cite their pass/block in the ratification comment.” Worth piloting at q7 before codifying.

Bucket 3 – Phase-2 design inputs (constraints, not architecture)

These are the audit-derived constraints any successor coordination shape (subagent-chain, agent-team, hybrid) should treat as load-bearing. The synthesis does not propose an architecture – only the boundary conditions an architecture must respect.

P1. Member load is dominated by a tight 4–5 chunk set

Across instances #2–#11 the chunks that drove decisions consistently were: (a) the PR body’s Convergence-state header, (b) the current `coo/proposed_<basename>.md` draft, (c) the last comment (= startup prompt), (d) the prior Instance #(N–1) report comment, (e) the prior 1–2 commit diffs on the committee branch. **Phase 2 should pre-fetch this chunk set deterministically per Member fire**; making each Member sub-agent re-pay the GitHub-MCP roundtrip tax is the wrong shape.

P2. CLAUDE.md boot is largely waste for Member rounds

- Identity files (charter / governance / preferences): 8 of 12 instances explicitly noted “skimmed and discarded, no decision-shape this round”.
- Mem0 recent-episodic search: 4 of 12 explicitly nil signal.
- Deferred-tools manifest + skills list: 5 of 12, “~95% never used”.
- SessionStart memo headers: 4 of 12, skim-and-discard.

These are session-identity artifacts, not per-round artifacts. **Phase 2 should decouple per-round context-load from session-identity-load.** A successor Member sub-agent type that inherits identity from a long-running Coordinator parent (rather than re-loading from CLAUDE.md every fire) would reclaim the bulk of the wasted token budget.

P3. Tool-friction is harness-environment, not coordination-shape

The four dominant friction patterns – `--repo` flag (9 instances), `--comments` overflow (6), `Read 25K` cap (5), branch mismatch (6) – would reappear unchanged in any subagent-chain or agent-team architecture. They are artifacts of running `gh` against a local proxy remote, harness-managed checkouts, and the `Read` tool’s hard token cap. **Phase 2 should assume these constraints persist** unless solved at the substrate layer; the protocol language should not pretend they don’t exist.

P4. Convergence-state in PR body is canonical shared-state

Every Member instance read the PR body’s Convergence-state header as the authoritative “where are we” surface; every Member updated it before posting. Instance #1’s report (extending the worked case) and the per-round handoff comments served as the **shared mutable state of the chain**, with the PR diff serving as the **shared mutable artifact**. **Phase 2 must preserve an explicit shared-state surface**

distinct from per-round handoff text. A subagent-chain without a written shared-state mechanism would lose the narrowing-catch convergence signal that worked across q3 and q6.

P5. Member sub-agents cannot themselves spawn sub-agents

The current `quorum_member_task.md` §“Do the work” calls this out explicitly: “You CANNOT spawn sub-agents (Explore, Plan, general-purpose) for parallel research — Claude’s sub-agents cannot themselves spawn sub-agents.” Phase 2 must respect this Anthropic-platform constraint: any “research delegation” capability for a Member-equivalent role must either live with the Coordinator (= Instance #1) or escape to a side-channel (separate routine fire, dispatched research session). The asymmetry is structural, not optional.

P6. Prior Instance reports grow in load-bearingness across rounds

Audits showed late-round Members consulting Instance #(N-1)’s report (and sometimes #(N-2)) as primary context — instance #11 cited #10’s handoff verbatim; #10 cited #9; #8 cited #7’s verification-claim discipline; #5 had to re-fetch the 2KB-truncated startup prompt mid-round. **Cumulative chain weight grows linearly with round count.** Any Phase-2 mechanism that replaces routine-fired Member sessions with subagent dispatch needs equivalent persistence; ephemeral sub-agent contexts won’t carry round-N report through to round-N+3.

P7. Coordinator/Member asymmetry is structural and worth preserving

The Coordinator’s audit spans commissioning → Instance #1 → arc observation → ratification → post-merge → retrospective. The Member audits cover one round each. Coordinator has sub-agent privileges (q3 #1 used four parallel Explores; q6 #1 used inline greps); Members do not. Coordinator holds long-running PR subscription and watches for chain-health anomalies; Members fire and exit. **Phase 2 should preserve this asymmetry or explicitly redesign it** — collapsing to a homogeneous architecture would lose the arc-observation primitive that caught the supernumerary instance #11 case.

P8. The forensic-audit primitive itself is reusable

Twelve instances filed usable audits — most live, none synthetic — within ~80 minutes of commissioning. The per-chunk schema (Source / Summary / Size / Use / Effect) plus three closing sections (Decision points / What was missing / Tool friction) was sufficient to run cross-instance synthesis without re-reading raw audit text at the synthesis layer. **Phase 2 instrumentation should preserve this primitive** — if subagent-chain replaces routine-fired chains, each sub-agent should still be able to file a structured audit on demand. The primitive’s reusability is itself a load-bearing property of the coordination substrate.

Cross-cutting observations (not actionable, but worth recording)

- **Standing-order text is more load-bearing than identity-file text.** Across audits, `committee_protocol.md` §3/§4/§5/§6/§7 sections were cited as decision drivers in nearly every Member round; `identity/charter.md`, `identity/governance.md`, `identity/preferences.md` were uniformly skim-and-discarded. The protocol earns its keep; the identity files do not for this work-shape. (Identity files are doing other work — session-identity posture and authority-bound — that the audits don’t measure.)
- **Pitfall list in the spawn guide pays off.** Instances #4, #5, #6, #7, #8, #9, #10, #11 all cited the spawn-guide Pitfalls list (#5 shallow-nitpick particularly) as a decision driver. Continuing to grow this list with explicit case-law (q1, q3, q6 lessons) is high-leverage.

- **The quorum label remained semantic-only post-q4.** No instance reported label-fire confusion, suggesting the q4 retrospective fix held. (Spawn guide §Pitfalls #9 codifies this.)
- **No instance reported confabulation this round.** The q3 instance #3 confabulation case-law held — instances surfaced uncertainty as drift-flags rather than confidently fabricating diagnoses. The verification-claim discipline that emerged in instances #6→#7→#8 is a direct continuation.

What this synthesis does NOT do

- **Does not redesign Phase 2 architecture.** P1–P8 above are constraints on the design space, not proposals.
- **Does not amend `coo/committee_protocol.md`.** The protocol is constitutional per its own §Scope; touching it requires a quorum pass. None of the recommendations above need protocol changes — they are operational.
- **Does not propose memo issuance.** Audits were the artifact; this retrospective is the synthesis. A memo would be appropriate only if Phase 1 amendments materially change the routine contract; they don't, they clarify.
- **Does not gate PR #176 merge.** PR #176 is shippable as-is — the Phase-1 amendments (A1–A5) are improvements, not blockers. If the BDFL prefers to merge #176 first and treat A1–A5 as a follow-up PR, that's clean.

Suggested follow-up

1. Apply A1–A5 as commits on `claude/quorum-routine-externalization` (PR #176) if still open; otherwise as a small follow-up PR.
2. File issues for O1–O7 with the appropriate `area:*` labels per the `tagging-taxonomy` skill. O1 (`gh --repo` discipline) is highest impact and not quorum-specific.
3. Carry P1–P8 into the Phase-2 design discussion as load-bearing constraints.
4. Close issue #183 once this synthesis is committed and the link-back comment is posted.

End of retrospective. Single-instance authority; commit + PR for BDFL read.

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) | — | — |

| # | Date | Topic | Implementing PR | Retrospective |
|---|------------|---|--------------------|---------------|
| 3 | 2026-04-24 | Memo system transition (per- file layout) | #89 | ... |