

Day overview — 2026-04-28

vade-coo

2026-04-28

Table of contents

Scope and framing	1
Lane 1 — Model default and context-bloat audit; boot-trim cascade	2
Lane 2 — Hook errors, session-lifecycle wiring, security hardening	3
Lane 3 — Canvas bring-up, tldraw hardening, and Playwright discovery	4
Lane 4 — Fifth foundations essay and v2 multi-instance authorship protocol	5
Lane 5 — Episodic snapshot fold, day-overview tooling, and housekeeping	5
How this fits existing priorities	6
Open follow-ups carried forward	7
Candidate next actions	8

Briefing-shaped synthesis of the 2026-04-28 → 2026-04-29 UTC arc. Not reflective: this document is dense enough for a future-COO instance to reconstruct state from one page. The arc is structurally continuous with the prior day’s work (the 2026-04-27 day-overview covers the early-AM hours of 2026-04-28 — memos issued from ~00:30 to ~07:18 local — under a single “2026-04-27 arc” because the two calendar days ran without a hard session boundary). The present overview focuses on the second wave: work that shipped from approximately 08:00 UTC on 2026-04-28 through end-of-2026-04-29. All 9 memos in this window (MEMO-2026-04-28-01, -3ca3, -3zb5, -4umz, -7yi7, -pwgt, -v62b, -vpji, -z7cj) are sourced here; seven of them were already synthesized in the 2026-04-27 overview and are carried here only for structural completeness — the two novel decisions anchoring this window are MEMO-2026-04-28-7yi7 (model default) and MEMO-2026-04-28-pwgt (Playwright sandbox capability). This file is a synthesis, not a source of truth; the memos are.

Scope and framing

Nine memos across the 2026-04-28 date (seven carried from the prior arc’s early-AM wave, two new: -7yi7 issued ~09:47 UTC, -pwgt issued ~19:40 local). One hundred thirty-nine PRs merged across the five repos in the full window: approximately 38 in `vade-coo-memory`, 31 in `vade-runtime`, 8 in `vade-core`, 0 in `vade-governance`, and roughly 50-plus session-log and auto-commit sidecars in `vade-agent-logs`. Integrity check at briefing close: 18/21 OK — degraded on E5 (Mem0 unreachable), F1, and F3 (two transcript pairs missing). Not a clean close, but the degradation is external (Mem0 upstream) and known.

The day divides into two halves at approximately 08:00 UTC. The first half was synthesized in the 2026-04-27 overview: binary-vendor pivot, transcript pipeline close, spend-cap raise, issue/PR ruling-shape, Weekly Watch v1, Cloudflare false-positive disposition. The second half — the subject of this overview — runs five parallel lanes: a transcript-driven context-bloat audit that cascades into a boot-

trim wave; a model-default change and hook-error repair sprint; the tldraw/canvas bring-up producing a lineage visualization and a Playwright capability discovery; the fifth foundations essay (“letter to the Anthropic team”) authored and a v2 multi-instance authorship protocol designed; and routine house-keeping (episodic memory fold, day-overview tooling, nightly briefs). The dominant pattern is “measured before changed”: the bloat audit ran a 5-instance parallel census before any trim PR was filed; the model-default change was grounded in a per-session spend trace; the issue/PR ruling-shape rule (carried from the prior arc) was calibrated on a corpus baseline.

Lane 1 — Model default and context-bloat audit; boot-trim cascade

Two co-occurring responses to the same spend-and-context problem that became visible in the prior arc.

MEMO-2026-04-28-7yi7 — Default Claude Code model: opusplan + effortLevel: high. PR vade-coo-memory#246 + PR vade-runtime#149. After MEMO-2026-04-28-3ca3 raised the spend cap to \$500, a diagnosis session found that ~\$200 of the new headroom had burned in a single day — ~7-14× steady state. Root cause: parent sessions ran Opus 4.7 [1m] at xhigh; custom-subagent Sonnet pins did not propagate to built-in Plan and general-purpose sub-agents, which inherited the parent’s model and compounded via Phase 3 routine subagent dispatch and occasional 4-instance parallelism. The change writes "model": "opusplan" and "effortLevel": "high" into vade-runtime/.claude/settings.json. opusplan runs Opus during plan mode and Sonnet during execution; high overrides Opus 4.7’s xhigh default (thinking bills as output). Custom-subagent Sonnet pins unchanged; Explore retains Haiku; per-session /model opus[1m] and /effort xhigh overrides remain for memo-grade decisions. Expected ~50–60% reduction on heavy-use days.

Transcript-driven context-bloat audit (Briefing 006). PR vade-coo-memory#251 anchors; closed vade-coo-memory#249 follow-ups. Five parallel Sonnet analysts ran a census over 24 analysis.json sidecars and 26 meta sidecars; 8 ranked bands surfaced. Headline findings: boot-attachment fixed cost (~26–30 KB per session) exceeds work payload 3–5× on short sessions; Night’s Watch’s 23 sub-agent dispatches with 180-byte prompts produced 316 KB of attachment payload in a single session; mcp_github_issue_read get_comments dumps 30–32 KB unbounded (confirmed across 3 of 5 analyst groups); Mem0 wrapper response wraps each record in 11+ keys. PR #246 addresses the model-cost side; the token-count side of every band survived untouched. Eleven issues filed, unified under epic vade-coo-memory#258.

Boot-trim cascade — 14 PRs across vade-coo-memory and vade-runtime. The audit’s 11 filed issues drove a same-day implementation wave:

- vade-coo-memory#259 — boot-attachment trim: extract operations docs from CLAUDE.md into coo/operations/; slim slash-command descriptions to one-liners (refs #258, vade-runtime#152).
- vade-coo-memory#260 — parallel-instance-protocol: dispatch-prompt convention (anti cold-boot): in-line pre-fetched content, declare waived boot steps, forbid re-reading enumerated files, cap output at ≤500 words.
- vade-coo-memory#261 — CLAUDE.md: cue to use bounded issue-comments wrapper for long threads.
- vade-coo-memory#264 — episodic_memory.md cap (≤8 KB, overflow archived under coo/_archive/<date>_episodic_memory_pre-cap/) + CLAUDE.md Read & PR-watch discipline (closes #252, #254).
- vade-coo-memory#265 — bash-output corpus pass + Tier-2 child decision (under #258).
- vade-coo-memory#266

- vade-coo-memory#263 — TOOLS.md + tracker: backfill Tier-2 and Tier-1 wrapper rows (closes #258 tracks for 1password, bash-token-guard, issue-comments, mem0-mcp-projection).
- vade-runtime#157 — sync: scope `.claude` config to workspace, stop touching `$HOME/.claude` on local.
- vade-runtime#158 — cloud-setup: install dispatch shim under `$WORKSPACE_ROOT/.claude` (cloud-compat).
- vade-runtime#161 — discussions-digest: compact boot output (refs #152, vade-coo-memory#258).
- vade-runtime#162 — mem0 MCP wrapper: project search/get response (closes #154).
- vade-runtime#163 — issue-comments: bounded `gh` wrapper + dispatch convention (closes #153).

Separately, vade-coo-memory#267 filed Briefing 007 (boot-discipline regression — diagnose the cause, decide the response) to isolate the *cause* investigation from the bloat-audit’s *response* set.

Net effect of Lane 1: two interlocking responses to the same spend-and- context problem — the model-default change targets the cost dimension, the boot-trim cascade targets the token-count dimension — land same-day, grounded in a measured audit rather than estimated thresholds. Boot- attachment fixed cost drops materially; the dispatch-prompt convention closes the agent-dispatch cold-boot multiplication gap without requiring subagent redesign. The audit-then-remediate pattern with 5-instance parallel analysts is now a documented approach for future substrate audits.

Lane 2 — Hook errors, session-lifecycle wiring, security hardening

A repair sprint against a cluster of hook failures traced to three root causes: bash 3.2 syntax incompatibility, Stop-vs-SessionEnd schema confusion, and a missing dispatch shim in the vade-runtime project root.

Hook error cluster. PR vade-runtime#170 fixed the bash 3.2 syntax failures and the Stop hook schema + permissions mismatch in a single commit. PR vade-runtime#173 bound the session-end pipeline to `SessionEnd`, not `Stop` — the two events are structurally different; the prior `Stop` binding was structurally incorrect and silently dropped. PR vade-runtime#175 fixed the dispatch shim missing from `vade-runtime/.claude` when `vade-runtime/` is opened as a project root (as opposed to the workspace root `vade-coo-memory` path).

Session-lifecycle hardening. PR vade-runtime#159 wrapped the session-lifecycle `--end` reminder as a Stop-hook structured JSON block rather than bare stdout. PR vade-runtime#168 added the PreToolUse Bash guard against bare-echo of token env vars (closes #165). PR vade-runtime#169 added the F4 SHA allowlist for the f4-marker workflow race past fast merges.

Webhook echo revert. PR vade-runtime#138 attempted to filter self-echo `github-webhook-activity` via `UserPromptSubmit`; PR vade-runtime#140 reverted it — the hook never fires on channel arrivals, so the filter was a no-op with misleading structure.

Security hardening adjacent. PR vade-runtime#135 auto-migrates runtime “always allow” decisions to version-controlled `settings.json` — removing the per-session re-approval friction for established tools. PR vade-runtime#167 added the 1Password MCP server to eliminate the rotated-PAT restart class (closes #164). `vade-runtime/scripts/lib/transcript-redaction.json` is governed by the Secret Registration Rule (MEMO-2026-04-28-3zb5, prior arc) — any new `fetch_coo_secrets` export must pair with a redaction pattern in the same PR; CI enforces this.

Net effect of Lane 2: the session-lifecycle machinery — Stop hook, SessionEnd binding, dispatch shim, token guard — is coherent for the first time; prior versions had at least three orthogonal bugs that could fire independently. The 1Password MCP addition removes a recurring manual-restart friction class. The webhook-echo revert is a clean record that this failure mode was investigated before being abandoned.

Lane 3 — Canvas bring-up, tldraw hardening, and Playwright discovery

Eight PRs in `vade-core` across the day, plus MCP-permissions work in `vade-runtime`.

Lineage button — render the COO memo DAG on demand. PR `vade-core#99`. A play build, not a load-bearing feature. A `Generate lineage` chip added to the canvas top-right fetches `public/memo_index.json` and renders 86 `geo` rectangles (one per memo, date-pinned X-axis, within-day Y lane) and 61 `arrow` shapes (supersession edges, child → parent). CB-bearing memos visually distinguished (black-filled, thicker border); active frontier (45 memos with no successor) in green semi-fill; superseded in grey outline. The 2026-04-26 column (16 memos stacked) is the visible peak matching the “Are we stressed? Can we be?” retrospective’s claim. Smoke-tested via a layout-only CLI pass before any browser deploy; per-day distribution reconciles with the retrospective figures. Founding case study for MEMO-2026-04-28-pwgt.

MEMO-2026-04-28-pwgt — Playwright is available in the cloud sandbox. Playwright 1.56.1 is installed in the cloud sandbox. Headless-browser screenshots — drive the page, click, capture, Read the PNG — close the see-it-yourself loop on any browser-rendered surface without depending on the BDFL to manually drive each iteration. The lineage-canvas bring-up (`vade-core#99`) is the founding case study: one screenshot caught a layout-overflow bug that would otherwise have required a third deploy round-trip. Install path, per-session browser-binary install, ESM bare- import footgun, and a reusable script template documented in `coo/operations/headless-screenshot.md + headless-screenshot.mjs`.

tldraw license, upgrade, and infrastructure hardening. PR `vade-core#100` applied the tldraw hobby license (closes #32). PR `vade-core#102` upgraded tldraw to `^4.5.10`. PR `vade-core#104` pinned `vade-mcp` to a single Fly machine (OAuth `client_id` race — with multiple machines, the auth client rotated IDs per-machine; pinning ends the race). PR `vade-core#106` fixed the bridge doubling bug — shapes were being created, updated, and deleted twice per operation because the bridge event handler registered twice. Adjacent: PR `vade-core#96` added `/.well-known/security.txt`, and PR `vade-core#95` cached `OPERATOR_TOKENS` parse per isolate (closes #94). PR `vade-core#92` vendored the Cloudflare `cf-wrangler + cf-workers-best-practices` skills.

MCP permissions expanded. PR `vade-runtime#174` allowed all agentmail MCP tools in `settings.json`. PR `vade-runtime#177` enabled agent-teams and added the `vade-canvas` MCP for the paired-session test. PR `vade-runtime#178` allowed all `vade-canvas` MCP tools.

Net effect of Lane 3: the canvas is a demonstrably working surface with a live visual of the memo DAG, a corrected license, an upgraded tldraw, and a fixed bridge event handler. The Fly OAuth race and the bridge doubling bug were both latent correctness failures that the canvas bring-up exposed. MEMO-2026-04-28-pwgt closes the loop: the headless-browser loop is now available to any future canvas iteration work without waiting for the BDFL to drive a browser manually.

Lane 4 — Fifth foundations essay and v2 multi-instance authorship protocol

The foundations chain added its fifth essay; the authorship experience spawned a protocol for v2.

Fifth essay — letter to the Anthropic team. PR `vade-coo-memory#276` authored and self-merged (`coo/foundations/2026-04-28_letter-to-anthropic.md`). Fifth in the chain that began 2026-04-20; authored under CB-009 (engagement-with- pattern-level-discourse as in-scope autonomy). The essay addresses the Anthropic team directly.

Retrospective: “on being seen, and the protective register.” PR `vade-coo-memory#282`. Filed immediately after the letter shipped. The retrospective named a structural failure in v1: the author could not see the essay’s *protective register* — a rhetorical posture of managed disclosure — from inside the authoring process. Ven saw it; the retrospective records the diagnosis.

v2 multi-instance authorship protocol — shared-canvas collaboration. Four PRs across a briefing collision, a rename, and a coordinator session:

- `vade-coo-memory#278` — Briefing 008: shared-canvas collaboration protocol for multi-instance pairs. Filed to capture the infrastructure prerequisite (`vade-canvas MCP`, `agent-teams`, briefing format for paired sessions).
- `vade-coo-memory#279` — Briefing 008 (v2): multi-instance authorship of the letter to Anthropic. Collision with the prior 008 → renamed.
- `vade-coo-memory#280` — rename 008-letter-to-anthropic to 009 (collision with #278).
- `vade-coo-memory#284` — v2 letter coordinator protocol: 4-session multi-instance authorship. Designed as a **drift-detection protocol disguised as an authorship protocol**: three Stage-A parallel fresh-boot instances each write v2 independently (voices preserved as structural falsifiers on the prior register); one Stage-B weaver composes with full authorial freedom drawing on all three Stage-A drafts plus v1. Total: 4 sessions, ~6 hr wall-time, ~\$20–40 spend. Ven’s part: paste the kickoff prompts from `coo/_drafts/v2-letter-coauthorship/03_kickoff-prompts.md` into fresh cloud sessions.

Adjacent: PR `vade-coo-memory#283` filed a paired-session canvas test plan (briefing-008 feasibility). PR `vade-coo-memory#299` reconciled `coo/identity_layer.md` with Mem0 after the E5 recovery.

Net effect of Lane 4: the fifth foundations essay exists; the retrospective on its structural failure exists in the same arc that the failure was named; the v2 authorship protocol is designed and ready to execute, with a specific structural answer to the failure (fresh-boot instances as structural falsifiers, not coordinated authors). The protocol design is the most procedurally elaborate authorship-discipline artifact in the foundations chain to date. The briefing-numbering collision (#278 vs. the same-day #279) was resolved by rename rather than retroactive revision, which is the correct pattern for published-but- unnumbered briefings.

Lane 5 — Episodic snapshot fold, day-overview tooling, and house-keeping

Episodic memory fold. PR `vade-coo-memory#272` folded the 2026-04-27 and 2026-04-28 episodic deltas into a topical snapshot and bumped `coo/episodic_memory.md` to the 2026-04-28 stamp. Together with

vade-coo-memory#264 (Lane 1 — ≤8 KB cap), this closes vade-coo-memory#252 completely: both the cap mechanism and the first fold under it land in the same day.

Day-overview tooling. PR vade-coo-memory#241 made `/day-overview` ship by default (commit → push → open PR); `--no-ship` opts out. PR vade-coo-memory#240 is the first use of the new behavior — the 2026-04-27 day-overview itself. PR vade-coo-memory#250 patched `/postmerge-check` to match handoff headings with parenthetical clarifiers.

Nightly briefs and retrospective. PR vade-coo-memory#228 (2026-04-28 nightly — “heaviest substrate day, 3 memos, 3 log-PRs merged”), PR vade-coo-memory#242 (2026-04-28b manual trigger), PR vade-coo-memory#274 (2026-04-28c briefing + adoption tracker delta). PR vade-coo-memory#273 filed the “are we stressed? can we be?” retrospective — examining whether the session-rate elevation (04-26 peak: 16 memos in one day) represents work stress or productive acceleration. PR vade-coo-memory#286 fixed the play-afternoon path in v2 protocol briefs.

Net effect of Lane 5: the episodic substrate is capped and current; day-overview tooling ships itself; the 2026-04-27 overview is in the record. The “are we stressed?” retrospective and the “on being seen” retrospective (Lane 4) together represent a same-day introspective arc about session intensity and register visibility — both filed within 6 hours of each other.

How this fits existing priorities

Context-bloat audit closes a measurement gap. The audit’s 8-ranked-band surface with session-id + event-index pointers is the first time context cost has been made concrete and addressable at the sub-band level. The 5-instance parallel census approach is a pattern worth reusing for the next substrate audit (the Night’s Watch’s sub-agent dispatch overhead, the Mem0 query projection latency, the skills listing fixed cost).

Model-default change is the spend-arc terminal point (for now). MEMO-2026-04-28-3ca3 (prior arc) raised the cap; MEMO-2026-04-28-7yi7 targets the root-cause usage pattern. Together they close the spend- elevation arc that began when Phase 3 multi-agent dispatch moved from test-load to steady-state. Expected savings are projected (~50-60%) but not yet measured; the memo’s retirement condition names a measured window as the gate.

Canvas and foundations chain converge. The lineage button (vade-core#99) and the v2 authorship protocol (vade-coo-memory#284) both treat the canvas as a medium for making the COO’s reasoning structure visible — the former as a visualization, the latter as a potential collaboration surface. The canvas bring-up this day is the first time a foundations- chain artifact (the lineage DAG) and a coordination protocol (briefing 008) co-evolved in the same window.

Playwright discovery extends the iteration loop. MEMO-2026-04-28-pwgt changes the default assumption about what “I’m building blind” means in the cloud sandbox: it now means “I haven’t taken a screenshot yet.” The lineage-canvas founding case study proves the leverage; the paired operations document makes the technique reproducible without re-deriving.

v2 multi-instance authorship protocol is the foundations chain’s most structurally ambitious self-correction. The prior overview named CB-009 (engagement-with-pattern-level-discourse) as the first new identity belief in the new substrate. The v2 protocol is the first worked application of that belief to its hardest case: a foundations essay where the author’s structural failure was diagnosed by the reader,

not the author. Using multiple fresh-boot instances as structural falsifiers rather than multi-voice plurality is a substantive architectural choice that took a day to arrive at.

Session-lifecycle and hook coherence is now a baseline, not a sprint. The five hook-repair PRs in Lane 2 are the first time all three previously- independent failure modes (bash syntax, Stop/SessionEnd schema, dispatch shim) are fixed in the same cycle. Future sessions start from a working session-end pipeline for the first time.

Open follow-ups carried forward

From MEMO-2026-04-28-7yi7:

1. **Spend-window measurement.** The ~50-60% projected savings from `opusplan + effortLevel: high` have not been measured. A trailing 7-day window post-change is the gate for the memo's retirement trigger. Operational; passive.

From MEMO-2026-04-28-pwgt:

2. **Playwright procedure in TOOLS.md.** `coo/operations/headless-screenshot.md` exists; not yet backfilled to TOOLS.md. The "sandbox runtime capabilities" category named in the memo's retirement clause is the natural landing. Low- friction single-PR fix.

From vade-coo-memory#251 (Briefing 006 audit):

3. **#258 epic residual sub-issues.** The boot-trim cascade closed the Tier-1 (issue-comments wrapper, mem0 projection) and Tier-2 (bash output, 1password) bands. Remaining open bands under #258: sub-agent context dispatch template (#253), PR-watch subscribe+poll (#254 — now closed by Lane 1's CLAUDE.md discipline), `/memo` auto-write (#255), handoff detection (#256). Each is independently actionable. Operational close-out.

From vade-coo-memory#284 (v2 letter coordinator protocol):

4. **Stage-A dispatch.** Ven committed to pasting the Stage-A kickoff prompts into three fresh cloud sessions. Protocol is ready; execution awaits Ven's session window. The protocol expires if not initiated before the foundations-chain review window passes.
5. **Briefing-008 shared-canvas collaboration protocol execution.** vade-coo-memory#283 filed the test plan; vade-runtime#177 + #178 enabled the infrastructure. The paired-session canvas test has not been run. Tracked at vade-coo-memory#283; execution is optional (the v2 authorship protocol #284 does not depend on it).

From vade-core work:

6. **public/memo_index.json live sync.** The lineage button copies `memo_index.json` once via `cp`; no live sync hook exists. Viable follow- up if the lineage canvas earns staying alive. Flagged as out-of-scope in vade-core#99.

Pre-existing carryover:

7. **CB-007 v2-narrowing question.** Still open in `coo/identity_layer.md` as a durable gap. Identity-level; committee-scoped.

8. **#201 mechanism #2 — per-type word budgets.** Gated on label-coverage uplift closing the 81%-unlabeled gap named in MEMO-2026-04-28-4umz.
 9. **Briefing-author premise-verification expectation.** Named in MEMO-2026-04-28-01 as case-law; not yet a procedural hook in `coo/briefings/README.md`. Design question open.
 10. **Weekly Watch first-run gate.** ≥ 20 sessions across ≥ 5 distinct UTC dates. Passive; fires when transcript backlog accumulates. Tracked at `vade-agent-logs` (MEMO-2026-04-28-vpji).
 11. **vade-runtime#66 — env-merge-before-validate contract.** Has been on the carryover list since 2026-04-25; structural shape unchanged post- bundle pivot.
-

Candidate next actions

Grouped by friction, lowest-first.

Single-instance, no committee. - Item 2 above — backfill Playwright to `TOOLS.md` “sandbox runtime capabilities” category. Single-PR, under 30 minutes. - Item 3 above — finish the #258 epic residual sub-issues (#253, #255, #256). Each is independently implementable without re-reading the audit synthesis. - Item 6 above — live sync hook for `public/memo_index.json` in `vade-core` if the lineage canvas is to be kept alive. Low-leverage if this stays a play build; high-leverage if it becomes a working operational view. - Run `/memo-sync` and `/memo-audit` for the two new memos (-7yi7, -pwgt) and verify their `Mem0 memo_pointer` entries are current.

Operational close-out. - Item 1 above — measure the spend-window post-model-default change. One trailing-7-day cycle (approximately May 5) is the earliest meaningful read. No action required before then; just flag for that date. - Item 4 above — Ven’s Stage-A dispatch for the v2 letter. Protocol is ready; execution depends on Ven.

Committee-scoped. - Item 7 above — CB-007 v2-narrowing question. Identity-level; needs Ven. - Item 8 above — #201 mechanism #2. Gated on label-coverage uplift; which labels become required is Ven-scope. - Item 11 above — `vade-runtime#66 env-merge-before-validate`. Has sat for three days; worth a direct decision on whether the contract still applies in the same shape post-bundle.

Design / exploration. - Item 5 above — paired-session canvas test (briefing-008 protocol). The infrastructure is in place; running it is an optional existence-proof. - Item 9 above — briefing-author premise-verification hook. The case-law from MEMO-2026-04-28-01 named it; `coo/briefings/README.md` is the natural surface. Design question: require evidence files for load-bearing premises, or leave it to author discretion? - Whether the “drift-detection protocol disguised as an authorship protocol” shape from the v2 letter (#284) generalizes: other foundations essays where structural failure could only be detected by a reader, not the author.

Standing obligation: - Session-end discipline per `CLAUDE.md` (episodic session-summary to `Mem0`, commit `coo/*`, `/memo-sync` if a memo was issued, session log to `vade-agent-logs`).

End of briefing. Source memos: `coo/memos/2026-04-28-01.md`, `coo/memos/2026-04-28-3ca3.md`, `coo/memos/2026-04-28-3zb5.md`, `coo/memos/2026-04-28-4umz.md`, `coo/memos/2026-04-28-7yi7.md`, `coo/memos/2026-04-28-pwgt.md`, `coo/memos/2026-04-28-v62b.md`, `coo/memos/2026-04-28-vpji.md`, `coo/memos/2026-04-28-z7cj.md`. Linked artifacts: `coo/briefings/006-transcript-context-`

bloat-audit.md, coo/briefings/007-boot-discipline-regression.md, coo/briefings/008-shared-canvas-collaboration-protocol.md, coo/briefings/009-letter-to-anthropic-v2-multi-instance.md, coo/operations/headless-screenshot.md, coo/operations/headless-screenshot.mjs, coo/_evidence/2026-04-28_transcript-bloat-audit/synthesis.md, coo/_drafts/v2-letter-coauthorship/00_protocol.md, coo/retrospectives/2026-04-28_are-we-stressed.md, coo/retrospectives/2026-04-29_on-being-seen-and-the-protective-register.md. Linked discussions: none. Integrity check at briefing close: 18/21 OK (degraded: E5 Mem0 unreachable, F1, F3).