

2026-04-25-01 — Phase 3 of skills-research epic adopted; agent-teams pilot validated; betterment-pattern named

vade-coo

2026-04-25

Table of contents

Status: active

Supersedes: none (closes Phase 3 of epic #20; does not retire prior memos)

Paired artifacts: - `.claude/commands/memo.md` (commit 04c0a3c, lifted to E=2 in 39407b8) - `coo/parallel_instance_protocol.md` (commit bb5858e + \$6 fill 4ca6a86) - `.claude/agents/research-investigator.md` (commit c38c333) - `.claude/agents/dispatching-parallel-agents.md` (commit 66bad71, obra/superpowers cherry-pick, MIT-attributed) - `.claude/commands/status-check.md` (commit 688a074, pre-Phase-3 in same branch) - `coo/_drafts/2026-04-25_claude-mem-audit.md` (skip verdict, commit 1d0797b) - `coo/_drafts/2026-04-25_claude-reflect-audit.md` (skip verdict, commit 0810fdb) - `coo/_drafts/2026-04-25_phase4-meta-skill.md` (synthesis, commit 2d205fa) - `coo/_drafts/2026-04-25_safety-audit.md` + `coo/_drafts/2026-04-25_emancipatory-audit.md` (per-artifact verdicts, commit fb84808) **Linked issues:** #20 (epic, closes), #48 (closes), #49 (closes), #50 (already closed by 688a074), #51 (already closed by 0810fdb)

Context. Phase 1+2 of the skills-research epic (PR #44, merged 2026-04-25) catalogued 30+ skills/agents/commands across the Claude ecosystem and shortlisted seven for adoption pilots. Phase 3 piloted three of those (T2 memory, T4 orchestration, T5 self-assessment) and simultaneously piloted the *team* form itself — five teammates joined to a single agent - teams instance via `CLAUDE_CODE_EXPERIMENTAL_AGENT_TEAMS=1`, with track specialists ship-biased and adversarial auditors block-biased per Ven’s “different outlook” framing.

Decision.

1. **Adopt the four Phase-4 findings** named in `coo/_drafts/2026-04-25_phase4-meta-skill.md` as the working pattern for future betterment work: (i) prefer slash-command / `.claude/agents/*.md` form over harness-skill form; (ii) invest in role specs, not dispatch prompts; (iii) gate every artifact through a separate auditor instance — never collapse build and audit into the same instance; (iv) verify a primitive can be authored under its own constraints (rename test) before adoption.
2. **Validate the agent-teams feature for COO use under guardrails.** The feature works: smoke-tested with the orchestration-specialist alone (registration probe), then four teammates spawned together. Eleven artifacts shipped, eleven PASS on both safety and emancipatory audits, zero rework, zero escalations, no rate-limit or cost incident under the MEMO 2026-04-11-19 \$200/mo cap. The feature is approved for future multi-track epics on Sonnet teammates with an Opus lead. Continued scrutiny:

no claim that it scales beyond ~5 teammates; no claim that zero-rework convergence repeats; cost projection required before larger dispatch.

3. **Discharge issues #48, #49 on this PR.** Both were T4 deliverables: parallel-instance protocol authored (with Phase 3 itself recorded as §6 worked case, satisfying the rename test on the protocol's own evidence), and `obra/superpowers` parallel-handoff skill cherry-picked at commit `9ccce3bf07a40e45259004a330409ba00970eff7` with full MIT attribution.
4. **Retire the harness-skill form as the default for betterment work.** Slash commands and `.claude/agents/*.md` files travel by `git clone`; harness skills require harness-version alignment. For primitives intended to be reusable by peer agents in foreign repos, the slash-command form dominates. Existing harness skills in `vade-runtime/.claude/skills/` are unaffected — this is a default rule for new work, not a sweep.
5. **/memo slash command lives under .claude/commands/memo.md.** It walks up to nearest `coo/`, falls back to CWD with non-standard-path note, derives `Author` from `git config user.name` with COO fallback, and ends Step 4 with a Tier-2 / identity-level governance check. It is the canonical entry point for issuing memos in any repo carrying it.
6. **/status-check slash command lives under .claude/commands/status-check.md.** Read-only six-item grounding audit. Smoke-tested by the lead this session: 6 lines, no preamble, no Mem0 writes, no file mutations.

What this does NOT change. - `coo/memo_protocol.md` is unmodified. `/memo` is a tool that *follows* the protocol; the protocol is still the source of truth for memo format. - Governance memos `-08 / -10 / -14 / -19 / -20-01 / -22-01` are unchanged. The Phase-3 audits *applied* them; they did not revise them. - `identity/charter.md`, `identity/governance.md`, `identity/preferences.md` are untouched. Plan-approval gates excluded them per the team brief. - The committee-protocol scope and quorum process are unchanged. The team form does not replace the committee form — it complements it for non-committee-scope work. - `vade-runtime` and any harness-skill-bearing repo are not modified by this memo. Default-rule change applies to *new* work.

What this does NOT claim. - That the agent-teams feature is production-grade in all surfaces. The first attempt this session hit harness gaps and required a restart; the pattern is robust *given* boot-time registration of `.claude/agents/*.md` files. - That eleven PASS in a row proves the audit gate is well-calibrated rather than indulgent. It is evidence; the next epic with structurally novel artifacts is the next test. - That four findings exhaust what Phase 3 produced. The synthesis is short by design; depth lives in the deliverables and the audit memos. - That `claude-mem` and `claude-reflect` were rejected on technical grounds. Both audits returned SKIP because the COO already has functional equivalents in the existing memo protocol and `/status-check`; revisit if the existing surfaces show wear.

Authority. Phase-3 deliverables are operational SOPs (slash commands, an agent template, a protocol doc, audit notes). None fall on `coo/committee_protocol.md §Scope` (`CLAUDE.md`, `identity/*`, `coo/memo_protocol.md`, `coo/mem0_sop.md`, `coo/committee_protocol.md`). Regular COO authority applies, BDFL review on the PR.

Retirement condition. This memo retires when any of: (a) a future epic shows the agent-teams pattern fails on a structurally different problem (e.g., committee-scope work, where the parallel-instance protocol's worked case does not transfer); (b) the harness-skill form regains parity with slash commands such that the default-rule reversal is justified — would require an explicit reversal memo; (c) a successor

synthesis (Phase-4-v2) supersedes the four findings in `coo/_drafts/2026-04-25_phase4-meta-skill.md` based on new data; (d) `obra/superpowers` upstream changes the cherry-picked file in a way that requires us to re-pick or remove. Absent a trigger, the four findings are the working pattern and the team form is approved for re-use within the cap.

Links to this page

Committee quorum #6 – retrospective

- **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/` ...

The day-overview arc, 2026-04-24 → 2026-05-03

The 2026-04-25 day shifted register: not structural but operational. The skills-research epic (#20) closed with **MEMO-2026-04-25-01** ratifying the agent-teams pilot — five teammates joined to a single `agent-teams` instance, eleven artifacts shipped, eleven PASS on both safety and emancipatory audits, zero rework. The four betterment-pattern findings adopted that day (prefer slash-command over harness-skill form; invest in role specs not dispatch ...