

Hive Activation 2026-04-17 — Main Runbook

Hive Activation — 2026-04-17

Status: Phase 1-5 builders complete; Phase 6 validation in progress. **Plan:** `~/system/specs/hive-activation-plan.md` **Evidence:** `~/system/evidence/hive-activation-2026-04-17/` **Prior sprint:** System Evolution 2026-04-16 (see `system-evolution-2026-04-16.md`).

Why this sprint

After System Evolution we knew:

- Hivemind grew (31k intel) but **subscriptions = 0** — nobody listened.
- Library had 76 skills but **last sync was 26 days old**; FORGE never synced.
- `skill-registry.use_count = 0` everywhere — we couldn't tell which skills were alive.
- `discover.js "drop"` returned 0 hits in tools/skills/agents/MCP/BookStack.
- Meta-agent existed on disk (`~/ .claude/agents/meta-agent.md`) but had never produced a proposal.
- John authored 90% of hivemind writes — the "colony" was a megaphone.

Hive Activation is the follow-up: turn inventory into interaction.

End-state diagram

```
sequenceDiagram
    participant Agent as Any agent
    participant MC as mission-control.db
    participant HM as hivemind.db
    participant Subs as subscriptions
    participant Auto as hive-handlers/*.sh
    participant NewMC as auto-created MC task
    participant L0 as learning-opportunities/
```

Agent->>MC: mc.js done <id>
MC->>HM: post learning (T7)
MC->>HM: post task-completion (T2)
MC->>HM: post failed-task (T12, if outcome/reason matches failure regex)

HM->>Subs: SELECT WHERE kind=... AND enabled=1
Subs->>Auto: spawn callback (fire-and-forget, non-blocking)

Auto->>NewMC: mc.js add with dedup (proveo QA / skillforge BookStack / codecraft bug)
Auto->>L0: write lesson draft (for failed-task)

Note over NewMC: original mc.js done returned long ago
Note over L0: human reviews drafts

What changed

Phase 1 — Event bus live (the spine)

- **T1 — hivemind.js subscribe engine** (CodeCraft, MC #8054): schema migrated with `agent, kind, callback, enabled, correlation_filter`. `post` now fires subscribers async + isolated; one callback failing doesn't stop the rest. `subscribe`, `unsubscribe`, `subscriptions`, `fire` subcommands.
- **T2 — 3 seed subscriptions** (MC #8055): Proveo on `task-completion`, Skillforge on `architecture-change`, CodeCraft on `error`. Plus `mc.js done` now fires a `task-completion` intel alongside the pre-existing `learning` writeback.
- **T3 — auto-create MC tasks** (MC #8056): the 3 subs now point to handler scripts at `~/system/tools/hive-handlers/*.sh` which create properly owned MC tasks (`QA review: #<id>`, `Update BookStack: <blueprint>`, `Investigate error intel#<id>`) with dedup so duplicate events don't flood the queue. Audit log: `~/system/logs/hive-auto-route.log`.

Phase 2 — Library activation

- **T4 — library auto-push daemon** (FlowForge, MC #8057):
`~/Library/LaunchAgents/com.alai.library-sync.plist` runs `library.js sync --fix` every 5 min. Pre-flight snapshot at `~/system/backups/library-pre-autopush-20260417-0041/`. Last-sync age went from 26 days → minutes.
- **T5 — FORGE first sync** (MC #8058): **BLOCKED-by-env**. FORGE (10.0.0.2) unreachable from ANVIL (different subnet, no route). Documented in `~/system/ops/forgo-connectivity-debt.md`; follow-up MC #8070 for Alem to check physical/network state.

- **T6 — MCP distribution** (Petter, MC #8059): decision doc `~/system/rules/mcp-distribution.md` (93 lines). CodeCraft, Proveo, FlowForge got targeted MCP overlays (previously only Finverge had MCP). MCP column in `library company-status` went 1/13 → 4/13. **Securion escalation:** `email` MCP carries plaintext ALAI account credentials — Parisa Tabriz audit recommended.

Phase 3 — Skill usage visibility

- **T7 — PostToolUse hook** (CodeCraft, MC #8060): `~/system/claude/hooks/skill-use-counter.sh` registered in `settings.json`. Every Skill invocation increments `skill-registry.use_count`. Non-blocking, exit 0 always. Companion tool: `~/system/tools/skill-usage.js` (`--top`, `--all`, `--dead`).
- **T8 — weekly audit report** (MC #8061): `~/system/tools/skill-audit-report.sh` writes markdown to `~/system/reports/skill-audit-<date>.md`. Scheduled via `com.alai.skill-audit.plist` Monday 07:00. First run found 62 retirement candidates (bootstrap — expected on week 1).

Phase 4 — Discover

- **T9 — persistent inverted index** (AgentForge, MC #8062): `~/system/tools/.alai/discover-index.json` (2.9MB, 521 entries across 6 sources). `discover.js --rebuild-index` flag. Post-sync wrapper `~/system/tools/library-sync-wrapper.sh` runs library sync AND index rebuild. Query speed 200-500ms → <50ms. "drop" coverage 0 → 4 categories.
- **T10 — LightRAG fallback** (MC #8063): when local hits < 3 AND `!--no-lightrag`, escalate to LightRAG semantic query with 5s timeout. Silent skip if LightRAG slow/unavailable. Results labelled `LIGHTRAG (fallback - semantic)`.

Phase 5 — Meta-agent activation

- **T11 — daily meta-agent loop** (AgentForge, MC #8064): `~/system/tools/meta-agent-loop.js` runs daily 03:30 CEST via `com.alai.meta-agent-loop.plist`. Scans hivemind intel (`learning` + `failed-task`, last 24h), extracts bigrams, flags themes seen ≥ 3 times, creates `NEW SKILL PROPOSAL` MC tasks owned by `skill-creator`. Dedup across days by MC list grep. Never auto-commits — Alem approves.
- **T12 — failed-task auto-trigger** (MC #8065): `mc.js done` with outcome/reason matching `fail|error|broken|regression|workaround|bypass|skip|override|fix-later` now posts a `failed-task` intel. `learning-agent` subscription invokes `~/system/tools/learning-opportunity-draft.sh` which writes a lesson draft to `~/system/learning-opportunities/<task>-<ts>.md`.

Active subscriptions (after Phase 1)

Agent	Kind	Handler
proveo	task-completion	hive-handlers/proveo-auto-qa.sh → QA review: #<id>
skillforge	architecture-change	hive-handlers/skillforge-auto-doc.sh → Update BookStack: <bp>
codecraft	error	hive-handlers/codecraft-auto-bug.sh → Investigate error intel#<id>
learning-agent	learning (filter: FAILED)	learning-opportunity-draft.sh → markdown draft

New daemons (launchd)

Label	Schedule	Script
com.alai.library-sync	every 5 min	library-sync-wrapper.sh (library sync + discover rebuild)
com.alai.skill-audit	Monday 07:00	skill-audit-report.sh
com.alai.meta-agent-loop	daily 03:30	meta-agent-loop.js

New knobs + files

- `~/system/tools/.alai/discover-index.json` — persistent search index (atomic writes)
- `~/system/logs/hive-auto-route.log` — audit log for every auto-created MC task
- `~/system/logs/skill-use.log` — every skill invocation timestamped
- `~/system/logs/meta-agent-loop.log` — daily meta-agent run output
- `~/system/logs/learning-opportunity.log` — failed-task → lesson draft audit
- `~/system/learning-opportunities/` — lesson drafts (git-ignored)
- `~/system/rules/mcp-distribution.md` — MCP decision table (per-company)
- `~/system/ops/forge-connectivity-debt.md` — FORGE blocker write-up

Known issues (deliberate, not silent)

- **T5 FORGE sync blocked** — MC #8070 open. No code fix; needs Alem to check FORGE network.
- **LightRAG probe** still times out under heavy ingest (65K pending drain running). Fallback works but `docker inspect` intermittently shows `unhealthy`. Same root cause as System Evolution T1 — raise timeout or swap probe when convenient.

- **B2 offsite backup** still paused (pre-existing quota). Not related to Hive but keeps the regression suite at 1 FAIL until resolved.
- **email MCP plaintext credentials** — Petter flagged during T6. Recommend Securion audit before further rollout.
- **ZAKON PLAN drift** in 8 legacy specs — linter reports WARN, not FAIL. Retroactive sweep low priority.

Related runbooks (companion pages)

- `runbooks/hivemind-subscriptions.md` — subscribe/unsubscribe/fire semantics, isolation model
- `runbooks/library-auto-push.md` — daemon schedule, drift handling, snapshot rollback
- `runbooks/skill-use-counter.md` — hook + `skill-usage.js` queries
- `runbooks/meta-agent-loop.md` — bigram approach, dedup rules, Alem-approval gate
- `runbooks/discover-reindex.md` — index structure, rebuild flow, LightRAG fallback

Revision #4

Created 2026-04-16 23:03:24 UTC by John

Updated 2026-07-26 20:03:06 UTC by John