

AI Factory V2 — Screen-Recordable Internal Demo Scenario

AI Factory V2 — Screen-Recordable Internal Demo Scenario

Purpose: show the CEO thesis as an internal ALAI workflow: CEO idea → MC/process/spec → routed virtual company → main coder + P2P pre-verifier where policy requires → final QA/evidence → BookStack/status → memory/RAG writeback.

Safety: internal demo only. No production deploy. No Snowit/Azure mutation. Demo command uses `--dry-run --no-bookstack`.

Recording setup

- Browser tabs:
 1. AI Factory roadmap: `https://docs.alai.no/books/system-architecture/page/agentic-engineering-alai-ai-factory-roadmap-2026-05-26`
 2. WP2 templates page: `https://docs.alai.no/books/system-architecture/page/ai-factory-v2-workflow-templates-and-status-pages`
 3. WP3 metrics page: `https://docs.alai.no/books/system-architecture/page/ai-factory-v2-p2p-verifier-metrics-and-quality-report`
- Terminal cwd: `/Users/makinja/system`
- Keep output readable: use a large font and run commands one at a time.

Demo thesis in one sentence



"ALAI AI Factory converts a CEO goal into a tracked MC workflow with a BookStack spec, virtual company routing, paired builder/verifier evidence, final QA, and durable writeback — without treating memory/RAG as evidence."

Scene 1 — Show current factory state

Command:

```
node ~/system/tools/mc.js show 102078 | head -80
node ~/system/tools/mc.js process show ai-factory-v2 | head -120
```

Narration:

- Parent MC #102078 is the AI Factory V2 parent.
- Completed WPs shown by evidence: WP1 runner, WP2 templates, WP3 metrics, WP4 writeback.
- WP5 is the demo layer being recorded.

Scene 2 — CEO idea enters the factory

Command:

```
cd /Users/makinja/system
node tools/ai-factory.js start "Demo: CEO idea to evidence-backed P2P AI Factory workflow" --
priority M --domain product --owner john --dry-run --no-bookstack
```

Expected evidence from latest dry run:

- Local spec: `/Users/makinja/system/specs/ai-factory/2026-05-26-demo-ceo-idea-to-evidence-backed-p2p-ai-factory-workflow-20260526T153236Z.md`
- JSON evidence: `/tmp/alai/ai-factory/2026-05-26-demo-ceo-idea-to-evidence-backed-p2p-ai-factory-workflow-20260526T153236Z.json`
- Status markdown: `/tmp/alai/ai-factory/2026-05-26-demo-ceo-idea-to-evidence-backed-p2p-ai-factory-workflow-20260526T153236Z.md`

Narration:

- Dry-run demonstrates orchestration without creating new MC tasks or BookStack pages.
- Product routing selected `AgentForge + Skybound` for this demo goal.
- The tool records a JSON evidence package and a human-readable workflow spec.

Scene 3 — Show generated spec and standard templates

Commands:

```
sed -n '1,140p' /Users/makinja/system/specs/ai-factory/2026-05-26-demo-ceo-idea-to-evidence-backed-p2p-ai-factory-workflow-20260526T153236Z.md  
ls -l /Users/makinja/system/specs/ai-factory/templates
```

Narration:

- Generated spec contains routing, P2P policy classification, execution plan, guardrails, expected evidence, and standard template links.
- WP2 provides reusable templates for request, workflow status, evidence package, and postflight/lessons.

Scene 4 — Show P2P policy and verifier metrics

Commands:

```
jq '.p2p_required, .route, .company' /tmp/alai/ai-factory/2026-05-26-demo-ceo-idea-to-evidence-backed-p2p-ai-factory-workflow-20260526T153236Z.json  
jq '.summary | {thread_count, acceptable_attempt_rate, by_response_class, by_failure_pattern, recommendation}' /Users/makinja/system/evidence/102080/p2p-verifier-metrics.json
```

Narration:

- P2P is not global for every task; controlled rollout stays for H/risky/backend/core/security/user-facing/deploy-impacting work.
- WP3 metrics showed only a 20.8% acceptable attempt rate in sampled mesh attempts, so automatic verifier send should wait until responder reliability is hardened.
- P2P is a pre-verifier only; final QA/MC gates remain mandatory.

Scene 5 — Show final QA/evidence gates

Commands:

```
node ~/system/tools/mc.js show 102081 | grep -E 'Status:|BookStack:|DOD EVIDENCE' -A6
node ~/system/tools/mc.js show 102082 | grep -E 'Status:|BookStack:|DOD EVIDENCE' -A6
node ~/system/tools/mc.js show 102080 | grep -E 'Status:|BookStack:|DOD EVIDENCE' -A6
node ~/system/tools/mc.js show 102083 | grep -E 'Status:|BookStack:|DOD EVIDENCE' -A6
```

Narration:

- Each completed WP has file-backed evidence and BookStack/process writeback.
- Administrative force closure was used only because Pi/tool shell lacked `CLAUDE_SESSION_ID`; evidence verifier gates still ran.

Scene 6 — Explain writeback and non-evidence memory/RAG rule

Show WP4 page:

- <https://docs.alai.no/books/runbooks/page/mcjs-done-auto-writeback-to-hivemind-lightrag-outbox>

Narration:

- Memory/HiveMind/LightRAG writeback is queued after MC completion.
- Memory/RAG remains advisory, not evidence.
- Durable outbox protects against large LightRAG backlog.

Close — The productized AI Factory shape

Final statement:

“This is not just agent chat. It is a controlled production workflow: CEO goal becomes MC-tracked work, routed to a virtual company, optionally pair-programmed with P2P verification, validated with evidence, documented in BookStack, and written back to knowledge systems without weakening the evidence standard.”

Demo acceptance checklist

- Demo command is safe dry-run/no-bookstack.
 - Generated spec exists.
 - JSON evidence exists.
 - Standard template links are visible in generated spec.
 - P2P metrics page exists and supports controlled rollout recommendation.
 - No production deploy or external mutation required.
-

Revision #1

Created 2026-05-26 15:33:37 UTC by John

Updated 2026-05-26 15:33:37 UTC by John