

Migration & Wiring

§6 Wiring

After `user.md` and `personality.md` are created, `~/claude/CLAUDE.md` must reference them in the Context Loading table.

Patch to add to the `## Context Loading` table in `~/claude/CLAUDE.md`:

```
```diff
| Before... | Load this |
+| Any session (user context) | `~/claude/user.md` – who Alem is, constraints, trust model |
+| Any session (agent identity) | `~/claude/personality.md` – who John is, hard constraints, de
 | Any task | `~/system/agents/specialist-mapping.json` |
```
```

The two new rows are inserted at the top of the table (load first, most general). No other changes.

Session boot sequence after wiring:

1. `~/claude/CLAUDE.md` loads (always, auto-loaded by Claude Code)
2. `user.md` loaded via Context Loading reference (Alem profile, constraints, trust)
3. `personality.md` loaded via Context Loading reference (John identity, hard constraints)
4. `bash ~/system/boot.sh` runs (tool-verified system state)
5. `node ~/system/tools/email-inbox.js pending` runs (CEO emails)

Wiring does NOT change the existing session start steps – it prepends the two context files to the start sequence.

§7 Migration Plan

Phase 1 – Create canonical shells (this MC's build task)

Create `~/claude/user.md` as an empty structured file with section headers only. Create `~/claude/personality.md` as an empty structured file with section headers only.

Deliverable: two skeleton files + 3-line CLAUDE.md patch. Estimated 30 minutes of build work.

Phase 2 – AI interview (live with Alem, ~20 minutes)

John runs the interview from §5 in a dedicated session. Alem answers all 13 questions. John captures the answers into `user.md` and `personality.md`.

Deliverable: populated `user.md` + `personality.md` (60-80 lines each), Alem-approved. This is the canonical state for the MC.

Phase 3 – Persona reference update (separate MC, low priority)

After Phase 2, update 78 persona files to reference `~/claude/personality.md` as the canonical source for persona information.

Estimated effort: automated `grep-and-append`, ~1 hour. Do NOT do this in the same MC as Phase 1 or 2.

§8 Acceptance Criteria

The Pillar #1 build MC is complete when all of the following are true:

AC1. `~/ .claude/user.md` exists and contains populated content in all §3 sections (verified by v

AC2. `~/ .claude/personality.md` exists and contains populated content in all §4 sections (verifi

AC3. `~/ .claude/CLAUDE.md` Context Loading table contains two new rows referencing user.md and p

AC4. Alem has reviewed and approved both files in a live session (verbal confirmation captured i

AC5. Session boot on a clean Claude Code launch loads both files without error (verified by runn

AC6. BookStack page published under Agentic OS shelf describing user.md and personality.md struc

AC7. No persona file content was auto-modified in this MC – confirmed by git diff showing zero c

§9 Risks and Non-Goals

Risks

R1: Interview capture drift. If John paraphrases Alem's answers rather than recording verbat

R2: CLAUDE.md already does this. CLAUDE.md contains identity information. Adding user.md and

R3: 131 persona files ignored. If Phase 3 never runs, the 131 persona files continue to infe

R4: Over-engineering the interview. 13 questions risks Alem losing patience at question 7. M

Explicit Non-Goals

- NOT changing behavior or content of any of the 131 persona files in this MC.
- NOT building any automation that reads user.md or personality.md programmatically.
- NOT generating user.md or personality.md content without the AI interview.
- NOT modifying any ZAKONS in CLAUDE.md as part of this work.
- NOT creating a new agent or skill for identity management.
- NOT making any changes to ~/system/agents/definitions/ in this MC.

§10 Cost Estimate

Design phase (this MC #99291): \$0 implementation cost. Spec writing only.

Build phase (separate MC):

- Phase 1 – shell files + CLAUDE.md patch: \$0 (file creation, no inference needed)

- Phase 2 – AI interview + composition: ~\$0.05-0.10 (one Claude session, Sonnet, ~20 min)
- Phase 3 – persona pointer update (separate MC): ~\$0.02 (automated grep-and-append)
- AC5 session boot test: ~\$0.01

****Total estimated build cost: < \$0.15****

No infrastructure changes. No new services. No API calls beyond normal session cost.

Revision #2

Created 2026-05-07 10:24:42 UTC by John

Updated 2026-06-07 20:01:23 UTC by John