

Skill: prompt-forge — Pre-flight team for H/BLOCKER tasks

Skill: prompt-forge — Pre-flight team for H/BLOCKER tasks

Skill path: `~/claude/skills/prompt-forge/`

Description

Pre-flight team meta-design for H/BLOCKER MC tasks. Spawns 5 specialist agents (petter-graff + devils-advocate + anthropic-chief-architect + openai-chief-architect + chip-huyen) in parallel. Synthesizer (petter-graff in 2nd role) combines into unified prompt with preserved `<disagreements>` XML field.

Output: `~/system/prompts/forged/<mc_id>.md`

Trigger: `/prompt-forge <mc_id>`

When to Use

Run prompt-forge when **ALL** of the following apply:

- Priority = **H**
- AND: **BLOCKER** label OR effort estimate \geq M OR 3+ acceptance criteria OR multi-component

Skip prompt-forge (go directly to /mehanic) when **ANY** of:

- Priority = M or L
- BLOCKER but trivial scope (single-file edit, single-line config, runbook execution)
- Recurring template task (deploy step, cert renewal, credential rotation)
- Effort estimate = S

Workflow

1. **Preconditions:** Verify MC task exists, status is open/in_progress, scope gate passes
2. **Parallel dispatch:** Spawn all 5 panelists simultaneously (Sonnet × 4, Opus × 1)
3. **Synthesis:** Petter Graff in synthesis role combines outputs into unified prompt
4. **Output format:** 8 sections including mandatory `<disagreements>` block
5. **Next step:** Feed forged prompt into `/mehanik`

Cost and Scope Gates

- **Hard cap:** \$1.00 per invocation
- **Daily budget:** \$5.00 (≈ 5 BLOCKER tasks/day maximum)
- **Typical cost:** \$0.34 warm cache / \$0.95 cold
- **Model tiering:** Sonnet × 4 panelists + Opus × 2 (Anthropic CA panelist + synthesis)

Output Format (8 sections)

The forged prompt at `~/system/prompts/forged/<mc_id>.md` contains:

1. OBJECTIVE — Single paragraph: what and why
2. CONTEXT FILES — Explicit list with absolute paths
3. DELIVERABLES — Numbered artifacts with acceptance signals
4. CONSTRAINTS — Hard limits (what NOT to do)
5. SUCCESS CRITERIA — Machine-verifiable checklist
6. AGENT ROUTING — Which specialist handles this
7. **<disagreements>** — Raw panelist conflicts, unsmoothed
8. FOOTER — Synthesis flag, cost estimate, next step

v0.1 Deferred (post-Akershus 2026-05-04)

- Eval harness: 20-task golden dataset replay
- Complexity classifier: automated scope gate
- cost-tracker.js integration
- Observability: per-panelist latency, token counts
- BFPR ship gate: ≥70% Builder First-Pass Rate over first 20 uses

References

- **Local SKILL.md:** `~/.claude/skills/prompt-forge/SKILL.md`
 - **Agent briefs:** `~/.claude/skills/prompt-forge/references/agent-briefs.md`
 - **Full spec:** `/tmp/std1-synthesis.md`
 - **Related:** ZAKON #25 (john-operating-system.md line 77)
-

Author: ALAI, 2026

Last updated: 2026-04-28

Revision #2

Created 2026-04-28 10:15:40 UTC by John

Updated 2026-05-31 20:06:40 UTC by John