

## 30 days. One team. One workflow. AI that actually moves delivery metrics.

The AI Context Sprint Lite takes one team and one workflow from AI tool access to measurable delivery improvement, with a reusable context catalog, lightweight guardrails, a working pilot of the selected workflow, and a workflow kit the team keeps after the sprint.

### WHEN THIS FITS

- **You know which workflow needs the help.** Either the AI Context Diagnostic surfaced it, or the pain is obvious from inside the team.
- **Your team already uses AI tools.** Copilot, Claude, ChatGPT, Cursor — adoption is fine, delivery hasn't moved.
- **You want a measurable before/after.** Cycle time, time to first review, revision rounds, defect catch. Metric defined in week 1, not week 4.
- **You can commit a small core team.** One lead, 1–2 practitioners, workflow access for 30 days. Without participation, this becomes a deck instead of a pilot.
- **You want to own the outputs.** Context catalog, prompts, and guardrails belong to your team after we leave.

### WHAT YOU BUILD IN 30 DAYS

- **Workflow baseline.** Concrete starting point on the metrics that matter. Baselines collected in week 1.
- **Minimum viable context catalog.** Reusable team context for the chosen workflow — decisions, conventions, gotchas — organized to be retrieved, not just stored.
- **Prompt and workflow kit.** Tested patterns the team uses every day — the prompts that worked on real tasks during the sprint.
- **Lightweight AI guardrails.** What AI can and cannot touch, review rules, data boundaries.
- **Pilot workflow delivery.** The chosen workflow runs end-to-end with the new system. We work alongside the team, not in a side lane.
- **Measured readout + expansion roadmap.** Before/after on the metrics defined in week 1, reviewer feedback, failure modes, and a roadmap for what to scale next.

### WHAT'S NOT INCLUDED

- Multi-team or org-wide rollout · Platform engineering or net-new tooling · Custom enterprise integrations · Security audit or compliance certification (security questionnaire packet available on request)

### SECURITY & LLM POSTURE

- **BYO-LLM can be scoped where supported.** Route inference through your Anthropic, OpenAI, Azure OpenAI, or AWS Bedrock account under your existing provider agreements. Default for early pilots is MemFlair-managed Anthropic Claude.
- **Security packet on request.** Data flow, processors, encryption, tenant isolation, retention, known gaps — covered before any data is shared. Email [matt@vetratek.ai](mailto:matt@vetratek.ai).

#### WHO BUYS THIS

CTOs, VPs of Engineering, Heads of AI, and product leaders ready to commit one team to 30 days of focused work on the workflow that matters most.

#### INVESTMENT

\$25,000–\$35,000  
30 days, design partner pricing. Standard \$40k–\$60k after first cohort.  
Diagnostic fee credits toward this engagement if signed within 30 days.

#### NEXT STEP

Email [matt@vetratek.ai](mailto:matt@vetratek.ai) to scope the sprint or start with the diagnostic.  
We'll confirm fit, agree on scope and metrics, and only commit if both sides see a real shot at measurable improvement.

**ABOUT VETRATEK** A workflow-improvement practice for engineering and product teams, helping teams move from individual AI tool usage to shared workflows, reusable context, and measurable delivery improvement.

**ABOUT MEMFLAIR** The workflow-context platform underlying this work — reusable team knowledge as AI-ready memory across tools, agents, and workflows. Design-partner co-build engagements available during the build window.