


A CTO'S GUIDE TO THE FOUR LEVELS

How AI-native is your engineering org?

A roadmap from prompt-level tools to autonomous, system-wide execution.



INSIDE

The four levels
Operating profile at each level
Six principles for moving up
Assess, plan, climb

Every CTO knows AI matters. Few have a clear path to scale it.

Ad hoc experimentation creates enthusiasm, but rarely changes the economics of an engineering organization. Without a framework, AI stays trapped in individual prompt windows. Engineers get personal speedups while review, release, and operations stay manual and disconnected.

The AI Transformation Model gives engineering leaders a staircase. It moves work systematically across four levels, from individual thought partnership to autonomous, system-wide execution.

ATtribution

Adapted from the AI Transformation Model published by Notion, developed with input from Ben Levick (Ramp) and Geoffrey Litt.

No framework is flawless and no two journeys are identical. The model provides a common vocabulary and the governance to move with conviction.

THE CORE THESIS

- [01] Move systematically from prompt-level assistant tools to shared, autonomous workflow agents.
- [02] Anchor every capability in a shared Company Brain of unified data, context, and explicit governance.
- [03] Deploy specialist agents with per-agent guardrails and deep integration into production systems.

AI-native is a spectrum, not a switch

Most engineering orgs are stalled at the surface: chat assistants and code autocomplete deliver modest personal time savings while delivery processes stay the same. Transformation begins when AI stops being a tool you query and becomes a colleague that executes work directly against live systems.

1. Diagnose

Map tools against concrete autonomy criteria.

2. Connect context

Expose verified internal knowledge through permissioned access.

3. Deploy agents


Move from text generators to active workflow participants.

4. Orchestrate

Coordinate many agents with human-in-the-loop governance.


Four levels, from thought partner to system.

Each level has a distinct goal and a distinct return. The jump that matters most for engineering is from Level 2 to Level 3, where AI stops waiting for a prompt and starts owning work.



Who starts the work

Initiation moves from people to events to the system



Where most engineering orgs are today.

L1 AI as a thought partner

Adoption is entirely user-driven. Engineers turn to web chat apps to explain an unfamiliar library error, generate a regex, or suggest test cases. The models are isolated from company records and know only their pretraining plus whatever an engineer pastes in. Individuals feel the relief immediately, but the impact is qualitative and nearly impossible to track.

Builds basic prompting fluency. Creates no durable advantage.

OPERATING PROFILE

- Interface
Web chat apps
- Context
Pasted text only
- Execution
Manual copy and paste
- Governance
Individual discretion

L2 AI as an assistant

Task-specific tools embed in daily software and use company context under data permissions. An IDE plugin indexes your repositories and follows shared architectural conventions. Search spans Slack, Drive, Figma, and internal wikis. IT and security teams connect data, govern access, enforce identity boundaries, and control whether proprietary code can be used for training.

Context-aware code suggestions

Drawn from private repositories and conventions.

Cross-app search

Semantic retrieval across tools and wikis.

Assisted escalation replies

Drafts grounded in docs and ticket history.

Meeting and inbox triage

Summaries, action items, priority by urgency.

MEASURED RETURN

3–5 h

Saved per engineer every week

Plus less rework, faster onboarding, and higher satisfaction.

THE CEILING

At Level 2, every workflow still needs a person to start every step. That is why delivery speed plateaus here.

The leap from Level 2 to Level 3 is agency.

At Level 3, specialist agents own recurring workflows and act across tools with human approval checkpoints. They trigger without a prompt, in response to operational events: a merged pull request, a failed test suite, an inbound customer ticket. Workflow owners set each agent's mandate, tool permissions, and criteria for human review. Agents investigate, draft complete resolutions, act across systems, and flag edge cases.

For a CTO, this is the first level where engineering output decouples from headcount. Capacity lost to process choreography comes back to the team.

10–40% Capacity reclaimed across core teams	3× Faster cycles from trigger to verified resolution	24/7 Continuous execution without waiting on people
---	--	---

What Level 3 looks like in engineering

[01] Autonomous coding agents Investigate error logs, find root causes in the codebase, write unit tests, and open draft pull requests.	[02] Escalation resolvers Diagnose multi-step issues, query backend databases, resolve escalations, and answer internal questions.	[03] Release and project managers Track delivery milestones, find blocked dependencies across teams, and generate changelogs in real time.
---	--	--

AI stops being an overlay. It becomes the operating system.

At Level 4, multi-agent systems run mission-critical workflows end to end and improve them over time. The engineering organization is rebuilt around agents that coordinate with each other.



What changes at Level 4

Multi-agent orchestration	Agents hand off tasks to each other and agree on outcomes without a person routing every step.
Self-improving loops	The system analyzes production failures, refines its own instructions, and updates test suites automatically.
Non-linear leverage	Customer volume and revenue grow without matching growth in engineering headcount.
Continuous governance	Policy is enforced by the platform at execution time, with deterministic guardrails and immutable audit logs.

Six principles for moving up the levels

01 You operate at several levels at once. Code review may sit at Level 3 while incident response is at Level 2.	02 Levels build on each other. They never replace each other.	03 Context is the foundation. Connected data unlocks every level above Level 1.
04 Not everyone must build. Typically 10–20% of people create the agents the whole team uses.	05 Start where you are. Connect chat use to company context before automating.	06 Use the model to align and assess. It is a shared language for the leadership team.

WHERE XHAWK FITS

XHawk takes engineering orgs from Level 2 to Level 4.

Most organizations stall at Level 2 because connecting models to production tools, live databases, and identity providers is technically hard and administratively fragile. XHawk combines specialist agents with the Company Brain, a real-time knowledge graph of your organization, and handles permissioned retrieval, event triggers, guardrails, and model orchestration across your existing stack.

Specialist agents

Pre-configured for product, engineering, and analytics, running around the clock with deep tool mastery.

Shared Company Brain

Unifies documentation, schemas, repositories, and historical business context.

Flexible event triggers

Cron schedules, GitHub webhooks, CI pipeline failures, or Slack messages.

Guardrails and deployment

Per-agent tool permissions. Frontier, open-weight, or fine-tuned models on public, private, or sovereign cloud.

Built for teams that require compliance, auditability, and deterministic control. XHawk is in the SOC 2 compliance observation period.

ACTION PLAN

Assess, plan, climb.

A three-phase blueprint for CTOs ready to move their organization from passive prompting to autonomous execution.

01

Assess

Audit current tool usage and context silos across engineering. Find which teams are at Level 1 and where Level 2 tools have real traction.

02

Plan

Select two or three high-friction, high-frequency workflows held up by manual coordination. Define acceptance criteria and human approval checkpoints.

03

Climb

Deploy XHawk specialist agents on your Company Brain. Measure reclaimed capacity, cycle time, and error reduction.

Start with the product

Explore it yourself

XHawk is usage based. Connect a repository, turn on your first agents, and pay only for the work they do.

Grow at your own pace

Add agents and workflows as each one proves its value. No long contracts or upfront commitments.

Talk to our team

Learn how teams like yours use XHawk to accelerate team velocity.

**Agents were the first step. AI
Employees are what comes next.
Build your AI-Native Operational System.**

xhawk.ai →

Start exploring XHawk today, or talk to our team about where agents can help most.