

A GUIDE FOR ENGINEERING LEADERS · OCTOBER 2026


# The AI Workforce for Outer Loop Automation

AI made coding faster. Automating the outer loop with specialist agents for engineering, product, and analytics makes the whole company faster.

## 92%

**"92% of U.S.-based developers are already using AI coding tools both in and outside of work."**

GitHub & Wakefield Research. **Coding is automated. The steps after it are not.**



## AI WORKFORCE

PR Reviews

SRE Agent

Data Analyst

CodeWiki

Product Manager

Chief of Staff

Brainstorming

Platform Engineer

Cost Optimization Agent


Synthetic Users Agent

# AI writes the code. The work after the code still runs on people.

Software engineering was never only coding. Teams have adopted AI at the keyboard, and coding agents now generate, build, and debug code in the inner loop. Once agents write the code, every step after it moves to the outer loop: verifying the change, operating it in production, and learning from what customers say. The outer loop is where the effort goes. Each of those steps needs company knowledge that no single person or chat window holds.

XHawk is the AI Workforce autonomy stack. It gives Product, Engineering, and Analytics teams specialist agents that run 24/7 in isolated cloud sandboxes, triggered by events, schedules, Slack, or other agents. Every agent works from one Company Brain: a permission-aware knowledge index over your code, documents, tickets, conversations, and data that lives inside your deployment and stays independent of any single model provider.

This paper describes the outer loop of engineering, the AI Workforce platform that automates it, the solutions it provides to each team, and the security and deployment model that lets enterprises run it on their own terms.



## COMPANY VELOCITY

**Individual AI plateaus.  
Velocity climbs once the  
outer loop is automated.**


[See page 05.](#)

## CONTENTS

|           |                                                        |    |
|-----------|--------------------------------------------------------|----|
| <b>01</b> | The Outer Loop of Engineering                          | 04 |
| <b>02</b> | The AI Workforce                                       | 07 |
| <b>03</b> | Solutions by Team: Engineering, Product, and Analytics | 09 |
| <b>04</b> | Security, Governance, and Deployment                   | 11 |
| <b>05</b> | Who XHawk Is For and Getting Started                   | 13 |
| <b>06</b> | Architecture Principles: Keep the Stack Separate       | 14 |

# Everything after the code is the outer loop.

The outer loop has changed because of AI. Code is now the fast part, and the role of the builder is changing with it. Builders still ship the product, but they also need to deeply understand what to build next. Every change has to be reviewed, tested, deployed, monitored, and measured against what customers asked for, using context spread across repos, tickets, runbooks, dashboards, and conversations.



## The new Outer-loop for AI-Native teams

| STEP                      | WHAT THE AGENT DOES                                             | AGENT                   |
|---------------------------|-----------------------------------------------------------------|-------------------------|
| Test                      | Checks work against the source                                  | Verification Agent      |
| PR reviews                | Reviews every PR with deep research                             | PR Reviews              |
| CodeWiki                  | Explains the code and keeps docs current                        | CodeWiki                |
| CI/CD                     | Puts gates to ensure only verified code ships to prod           | Platform Engineer       |
| SRE/triage                | Runs runbooks and finds root cause                              | SRE Agent               |
| Cost optimization         | Optimizes token and cloud spend continuously, not once a year   | Cost Optimization Agent |
| Data analysis             | Answers data questions in plain English                         | Data Analyst            |
| Customer feedback         | Ranks customer pain themes                                      | Product Manager         |
| Requirements              | Drafts specs from past decisions                                | Product Manager         |
| User testing and feedback | Simulates users to test flows and give feedback on new features | Synthetic Users Agent   |

# Company velocity moves only when the outer loop does.

Individual AI makes one engineer faster, and the gains stay with that engineer. Delivery velocity across the company plateaus because review, release, deployment, and operations still run at human speed. The curve bends upward only when the lifecycle itself is redesigned around agents.




Figure 1. Individual AI versus institutional AI across the software lifecycle.

## How XHawk delivers each stage

### 1. Share context

The Company Brain codifies institutional knowledge so every agent starts informed.

### 2. Automate the outer loop

Specialist agents run reviews, tests, CI/CD, triage, and analysis around the clock.

### 3. Change how teams work

Automations and cost tracking measure the organization on delivery, not output.

# Agents that run the outer loop

## Verification Agent

Checks the work before anyone acts on it.

"Review this PR against our conventions before a human sees it."

**Standards at scale** across every agent's output

**Catch the silent break** when a definition shifts underneath a metric

**Citations on every claim** so readers can check the work

GitHub · GitLab · Snowflake · Looker · Datadog

## Platform Engineer

Works the backlog while the team sleeps.

"Take ENG-412, ENG-418 and ENG-425 from the backlog and open pull requests."

**Ticket to implementation** with repo context and conventions loaded

**Automated code reviews** on every pull request, not only the ones a senior engineer had time for

**Continuous docs** and codebase-wide migrations, opened for review

GitHub · GitLab · Bitbucket · Jira · Linear · CircleCI

## SRE Agent

First responder that has already read the runbook.

"What changed before the 02:14 alert fired?"

**Always on.** Watches alerts, executes the runbook, keeps the comms thread updated

**Root cause analysis** that correlates the deploy, the alert, and the preceding change

**Postmortems and alert triage** that separate real signal from noise

PagerDuty · Datadog · Sentry · Grafana · Slack

## CodeWiki

Answers what the codebase actually does today.

"Walk me through what happens when a webhook arrives."

**Onboard without a buddy** from current code, not a stale README


**Find where it lives.** Point to the file and function that own a behavior

**Living documentation** that describes what actually shipped

GitHub · GitLab · Notion · Confluence


# Separate knowledge from intelligence.

AI needs more than intelligence. It needs context. XHawk is built in layers so that each can evolve on its own: specialist agents that do the work, a Company Brain that holds what your organization knows, and infrastructure that supplies models, connectors, governance, and the cloud they run in. Page 14 describes why the layers stay separate. Your agents work together. You stay in control.



## How an agent completes a task

Each task runs through a self-correcting loop that calls a language model only when a step needs it. Results loop back to retrieval until the work is done.



# From chat assistants to an always-on workforce

A background agent is an autonomous AI worker that executes tasks without requiring a human at the keyboard. Unlike interactive coding assistants, background agents are triggered by events, run on schedules, or coordinate in fleets.

| DIMENSION          | CHAT ASSISTANT       | XHAWK AI WORKFORCE                           |
|--------------------|----------------------|----------------------------------------------|
| Where it runs      | Your local machine   | Cloud sandbox                                |
| How triggered      | You type a prompt    | Event, schedule, or fleet                    |
| Scope              | Depends on the user  | Entire company knowledge, repos, and systems |
| Availability       | When you're active   | 24/7 continuous                              |
| Tracking           | Limited visibility   | Per-task execution and CFO-level tracking    |
| Organization model | Human-driven tooling | AI-native operational system                 |
| Model strategy     | Usually one model    | Best model selected per task                 |

## Platform capabilities

### Agents that work as a team

In auto-pilot mode, agents call each other to finish long tasks.

### Action across 100+ tools

Multiplayer agents collaborate with your team in Slack, Microsoft Teams, and the XHawk dashboard.

### Automations and batches

Trigger by schedule, GitHub events, CI failures, Slack, or webhooks. Queue work and finish the backlog overnight.

### Agents that remember

Carry what matters between conversations. Every action is added to the Company Brain.

### Scoped access per agent

Control which repos, MCP servers, and secrets each agent can reach.

### The right model for every task

One interface routes each task by complexity, latency, and cost.

# One platform, multiple teams. Compounding value.

Every team reads from the same Company Brain, so the context one team adds makes every other team's agents smarter. Just @XHawk the work, and agents take it from there.

---

## Engineering

PR Reviews · Platform Engineer · SRE Agent · CodeWiki · Verification Agent

Automate the outer loop. Agents handle code reviews, documentation, CI/CD, backlog-to-PRs, testing, and incident response, so engineers spend their time on design and judgment.

### Backlog to pull requests

Tickets become reviewed PRs overnight.

### Escalations in Slack

Debugging context brought to where engineers coordinate.

### Recurring page to ticket

Evidence attached, instead of another silent acknowledgement.

### Docs that stop rotting

API and user docs updated from merged changes.

---

## Product

Product Manager · Brainstorming · Company Brain

Be relentlessly customer and data obsessed. Agents turn scattered signal from tickets, calls, and threads into decisions you can defend.

### Spec drafting

Requirements grounded in your docs, tickets, and past decisions.

### Customer feedback synthesis

A ranked list of what customers keep asking for.

### Launch readiness

A release checked against its specs, docs, and open issues.

### Pressure-test a hypothesis

Argue the other side of a bet with your own data.

# Data and Analytics

Data Analyst · Verification Agent · Company Brain

Every teammate gets an analyst, in plain English. The agent writes the query, runs it against your warehouse, and explains what the answer means. Analysts spend their time on modeling and definitions instead of ad-hoc requests.

## Self-service analytics

Any teammate can ask a data question in plain English.

## Report drafting

Raw queries turned into narratives someone can act on.

## Numbers checked at the source

Every figure traced to the warehouse before it reaches a slide.

## Metrics standardization

Codify definitions so the whole company speaks one number.

## Morning insights

What moved more than 10% and why, delivered on schedule.

## Grounded in your stack

Snowflake, Databricks, PostgreSQL, and Looker.

# Executives

Chief of Staff · Company Brain

Company-wide visibility without the assembly work: answers about the business from live data, cross-team rollups that are current when opened, and the true cost of agent work by team and task.

## SHARED FOUNDATION

# The Company Brain: the shared memory every other agent works from.

## Grounded answers

Every answer tied to your data, context, and knowledge.

## Institutional memory

The context behind past decisions and trade-offs.

## One source across teams

No private versions of the truth.

## Onboarding context

New hires brought up to speed on any project.

# Built from the ground up for the most secure workloads.

Every session runs in an isolated sandbox, and your data is never sold, shared, or used to train models. Security is how the product is designed, not a feature added afterwards.

**Never trained on your data or code**

Your data and codebase are used to do the work you asked for, and nothing else.

**Isolated by default**

Every agent runs in its own sandbox, walled off from other sessions and other users.

**Encrypted end to end**

Data and credentials encrypted in transit and at rest. Bring your own keys.

**SOC 2 compliant**

XHawk is SOC 2 compliant and builds every component to its standards.

## Governance and audit

|                                   |                                                                                                                                                                  |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Permission-aware retrieval</b> | Answers are retrieved under the permissions of the person or agent asking, inherited from the source system. Nobody sees a document they could not already open. |
| <b>Scoped agent access</b>        | Each agent is limited to the repositories, tools, and secrets you grant it.                                                                                      |
| <b>Full traceability</b>          | Every decision, action, artifact, and agent interaction is traced and audited.                                                                                   |
| <b>Measurable AI cost</b>         | Token usage, infrastructure cost, and execution efficiency tracked by agent, workflow, and feature.                                                              |
| <b>Model independence</b>         | The index sits behind the model boundary. Switching providers is a configuration change, not a migration. Bring your own key.                                    |

## Deploy the way your security team requires.

Choose a deployment model based on your security and compliance needs. Every option runs the same platform and the same agents.

### FASTEST START

#### Cloud

Fully managed, secure, and compliant. The fastest way to get started.

### BEST OF BOTH

#### Hybrid

XHawk manages the control plane while your data stays in your infrastructure.

### MAXIMUM CONTROL

#### Cloud-Prem

The full platform runs in your AWS, GCP, or Azure account. Your data never leaves your VPC.

## Connect the tools your work already lives in

Slack, Microsoft Teams, GitHub, Jira, Zendesk, Notion, Google Drive, and 100+ more are indexed into the Company Brain, permission-aware and portable across every model you run. For internal systems with no off-the-shelf connector, there are three ways in.

### MCP servers

Point the brain at any Model Context Protocol server; its tools become available to every agent.

### Indexing API

Push documents from internal systems with their permissions attached.

### Forward Deployed Engineers

XHawk engineers build connectors and custom automations against your real workloads.

## Multi-model by design

Route work across Claude, GPT, Gemini, Grok, GLM, DeepSeek, Kimi, and open-weight or private models without vendor lock-in. When the market moves, change the model; your index, permissions, and history carry over untouched.

## Helping teams analyze, operate, and build better.

### FOR FOUNDERS

#### Scale output, not headcount.

Give every employee an AI workforce that builds, analyzes, automates, and improves the business.

### FOR CFOS

#### Make AI costs measurable.

Track token usage, infrastructure costs, and execution efficiency by agent, workflow, and feature.

### FOR CTOS

#### Automate the Outer Loop of Engineering.

AI agents handle code reviews, documentation, CI/CD, backlog-to-PRs, testing, and other repetitive engineering work.

### FOR VP PRODUCT

#### Be relentlessly customer and data obsessed.

AI Data Analysts deliver instant answers from trusted product metrics, customer insights, and experiments.


## GETTING STARTED

## Put your AI Workforce to work in minutes. Drive ROI in six weeks or less.

- 01 Connect your data, code, and business apps.** Agents get secure access in isolated sandboxes.
- 02 Configure agents.** Choose pre-built agents or define custom behaviors.
- 03 Create work once.** Trigger by schedule, GitHub events, CI failures, Slack, or webhooks.
- 04 Operate with AI as the system.** Agents handle the mechanical work; people focus on what only humans can do.
- 05 Custom automation.** Forward Deployed Engineers automate your company-specific requirements.

# Keep the layers separate.

Models improve every few months and infrastructure choices change with regulation and cost. Company knowledge compounds for years. The goal for every company should be an AI Workforce stack whose layers are separate and pluggable, so that no layer depends on a single model or on a single location where the agents run.



## 04 AI Workforce

Specialist agents you add, remove, or customize without touching the layers below.

## 03 Company Brain

Your knowledge index and permissions. It stays yours, inside your deployment.

## 02 Multi-Model Inference

Frontier, open-weight, or fine-tuned models, routed per task and swapped by configuration.

## 01 Public, Private or Sovereign Cloud

Run where your security and compliance teams require, and move when they change.

## What separation buys you

### Pluggable

Replace any layer when something better arrives. The rest of the stack keeps running.

### Model independent

Change providers without re-indexing, re-permissioning, or losing history.

### Location independent

Cloud, hybrid, or your own VPC. The same agents run in each.



**Put your AI workforce  
to work today.**

© 2026 XHawk. All rights reserved.

[xhawk.ai](https://xhawk.ai)