

# A new foundation for code.

A graph-native code repository for people and AI agents.


Troy Fortin / Founder / Firelock, LLC / September 2026

# The same map gets rebuilt.

Change a shop's discount rule. The cart, checkout and receipt may share the same code.

Find the code → Trace connections → Review the change

The next agent and reviewer can repeat that investigation.




Illustrative software relationships. Not a Kin result.

# The connections belong with the code.

Kin keeps code, the connections it finds and their history together.

When the code changes, you can look back at the connections Kin recorded with it.

Build from a persistent record.  
Not only the current session.



Illustrative model. Supported tools still use files.

Repository model. Kin can miss connections; compilers, tests and review still matter.

# AI adoption is widespread. Confidence still takes work.

## ADOPTION

90%

of surveyed technology  
professionals use AI at work.

[DORA 2025 / 2026 analysis](#)

## CONFIDENCE

46%

of respondents distrust the  
accuracy of AI output.

[Stack Overflow / 2025](#)

## REVIEW BURDEN

85%

agree AI shifted the bottleneck  
to reviewing and validating  
code.

[GitLab / The Harris Poll / 2026](#)

Generating a change is not the same as establishing confidence in it.

Different surveys, questions and populations. These indicators are not a combined statistic. Sources and scope in A5.

# Ask the repository what connects.

8

## Recorded callers.

One React helper connects to  
input, select and text-area  
code.

Kin v0.7.3  
React 9b7a0d40

Capture and limits [↗](#)

Single-commit snapshot; call coverage partial. Full provenance and capture limitations in A1.

Recorded React query

Prepared graph

kin refs getToStringValue --kind calls

ReactDOMInput.js

hydrateInput

initInput

updateInput

ReactDOMSelect.js

hydrateSelect

updateOptions

ReactDOMTextarea.js

hydrateTextarea


initTextarea

updateTextarea

Recorded callers, regrouped by file. Not a new live run.


# Context beside Git. Or part of the repository.

GIT + SEPARATE CONTEXT



People and agents query the derived record.

KIN'S REPOSITORY MODEL



People and agents query repository state.

The bet: a shared software record that stays with the revision being changed.

Indexes can be persistent and revision-aware. Selected models, not a performance ranking. Product examples and official sources in A5.

# Start where context already costs time.

INITIAL TARGET HYPOTHESIS

AI-heavy teams.  
20 to 200 engineers.

|                |                                                  |
|----------------|--------------------------------------------------|
| Buyer          | VP Engineering or platform lead.                 |
| Trigger        | Agent-driven changes outpace review capacity.    |
| First workflow | Shared-code changes in one supported repository. |

GitHub Octoverse 2025. Initial target segment for evaluation. Bottom-up sizing and pricing tests in A3.

THE SCALE OF THE WORKFLOW

43.2M

pull requests merged per month on  
GitHub in 2025.

Workflow volume, not Kin's market size.

One engineer and their agent.  
One recurring change.  
Beside the existing Git workflow.

# Earn the developer. Become the team's shared record.

AVAILABLE TODAY

Public alpha

## Kin

Local repository and graph queries.  
CLI, editor and MCP access.  
No hosted account required.

KINLAB / IN DEVELOPMENT

Proposed paid scope

## KinLab

Shared repository state.  
Team access and permissions.  
Persistent change evidence.

PRICING EXPERIMENT

Active-developer subscription or team fee plus usage.  
Test willingness to pay through fixed-scope paid evaluations.

Local Kin is Apache-2.0. Hosted paid scope and pricing are design-partner hypotheses, not available self-service features.

# The alpha is built. Now earn repeat use.


6 months

## Founder-built public alpha.

Native repository operations.  
Git import and export.  
CLI, editor and agent access.

Founder-confirmed September 8 snapshot: no revenue or signed design partners. Installation completion and retention have not been established.

EARLY CONVERSATIONS / SEPTEMBER 8



MEASURE NEXT

Activate

First useful answer

Return

Four-week repeat use

Pay

Team commitment

# Built from firsthand experience.

## Troy Fortin

Machine learning engineer building production generative-AI systems. Previously at The Home Depot.

Architected and shipped Kin's public alpha solo.

Self-funded. Full-time on Kin at financing.

### FIRST TECHNICAL HIRES / RISK REMOVED

---

01 Repository / runtime  
Production reliability

---

02 Developer / agent integrations  
Adoption and repeat use

---

03 Cloud / platform  
Shared deployment and security

Founder materials, September 7, 2026. Firelock, LLC. Hiring is staged against product readiness and customer pull.

# Capital should buy evidence. Not just features.

---

01 / RELIABLE USE

## Qualify the workflow.

5 to 10 external repositories in daily-use qualification.

4 weeks of continuous Kin-on-Kin use.

Track activation, coverage and setup cost.

---

02 / RETAINED TEAMS

## Earn recurring work.

5 to 7 design partners.

Measure four-week return and repeated use on real changes.

Expand from one workflow to another.

---

03 / PAID EVIDENCE

## Prove the commitment.

2 to 3 paid pilots.

At least 2 referenceable customers.

Independent reproducible evidence.

Track conversion, expansion and margin.

Operating targets for the 24-month plan, not achieved results. The repository count is a qualification cohort, not an acquisition ceiling.

# AI changed who writes code. Kin changes what they build on.

## RAISING SEED FINANCING

### From public alpha to daily team use.

A 24-month plan for repository reliability, agent adoption and shared KinLab workflows.

## WHAT THIS ENABLES

Retained external use.  
Paid team workflows.  
Independent evidence.


[Discuss Kin ↗](#)

Firelock, LLC

Operating plan and capital allocation in A4.

# The result is inspectable. So are its limits.

Eight recorded callers of `getToStringValue`, grouped by file.



v0.7.3 / React 9b7a0d40. Single-commit snapshot. Full history was refused on that build.

Reported JavaScript call coverage: 13,819 of 36,234 / 38% / partial.

Saved historical command output. Re-typeset, not rerun. It does not establish speed, completeness or end-to-end correctness.

# Prove the workflow. Not only the query.

Compare Kin and an indexed Git workflow on the same pinned change.

---

## HOLD CONSTANT

### One controlled task.

Repository revision, change request, model, context budget and tests.

Include graph setup and maintenance, not only warm queries.

Coverage, import scale and workflow support remain alpha constraints. Compilers, tests and review remain necessary.

Proposed evaluation, not a completed experiment. Live proof scope separates older retrieval results from the shipping release; no current accuracy or savings claim is added.

---

## MEASURE THE OUTCOME

### Useful work, end to end.

Time and tool calls to establish impact. Missing connections. Review rework. Finished-change quality.

Repeat on unfamiliar supported repositories.

# Size the buyer. Test what they pay for.

Qualified  
accounts

Count the initial segment

Annual seat-based opportunity model. A team-plus-usage plan needs a separate account model.

TEST 01

Active developer.

Subscription tied to actual use.

×

Active users  
per account

Measure participation

TEST 02

Team + usage.

Shared platform fee with usage.

×

Annual price  
per user

Test paid commitment


TEST 03

Paid evaluation.

Fixed scope and an explicit fee.

Pricing and validation experiments under active evaluation. Account metrics and commercial terms are tested with design partners.

# Fund the team. Stage the commitments.



## 24-month base plan.

Founder plus three technical hires. No uncontracted revenue assumed.

## A smaller close changes the scope.

The separate 18-month plan is not the same milestones for less money.

Go-to-market hiring is gated on at least two paid pilots and repeatable onboarding.

Planned allocation and operating assumptions, company-reported.

# Clear sources. Bounded claims.

## Company and evidence

Company-reported, September 7, 2026: product model, saved React output, outreach, founder and operating targets.

Company-reported, September 8, 2026: narrowed buyer, pricing experiments and evaluation plan.

[Kin source](#)

[Current public proof scope](#)

The React result is historical. Outreach numbers reconfirmed September 8. Conversations with developers, potential design partners and investors have just begun.

## Industry signals

[DORA / 2025 and 2026](#)  
90% AI use at work.

[Stack Overflow / 2025](#)  
46% distrust AI accuracy.

[GitLab / Harris Poll / 2026](#)  
85% review-bottleneck agreement; 1,528 respondents, six countries.

[GitHub Octoverse / 2025](#)  
43.2M merged PRs per month.

## Selected alternatives

[GitLab Orbit / beta](#)  
[Sourcegraph / SCIP indexes](#)  
[Augment / Context Engine](#)  
[Greptile / graph context](#)  
[GitHub / Copilot review](#)

Official product descriptions checked September 8, 2026. Not a feature ranking or a measured comparison.

Indexing does not mean ephemeral context or absence of revision support.