

AI CHANGE ASSURANCE

2nd Engineer

Ship at AI speed. Keep engineering control.

Independent assurance for AI-authored software change.
Separate review, exact-state evidence and human approval.

Selective by project

Independent review

Exact-change evidence

Owner controlled


THE MANAGEMENT PROBLEM

AI creates more code. Your review capacity did not.



THE GATES

Three coloured routes. One accountable decision.



VALUE FOR ENGINEERING LEADERS

Scale output without scaling ambiguity.



OPERATING MODEL

Fits around the AI coding workflow your team already uses.



Code volume can scale. Review discipline can scale with it.

2nd Engineer makes the second opinion selective, independent, reproducible and answerable to a person.

AI CHANGE ASSURANCE

Built as control infrastructure, not another comment bot.

The reviewer is only one component. The product value is the governed chain around it: owner authority, selective routing, independent context, immutable evidence, enforced outcomes and a protected control path.

THE CONTROL STACK

Six layers that preserve accountability

1

Owner authority
Durable OFF, MANUAL and AUTO modes. Only direct owner action can resume automatic collaboration.

2

Deterministic risk router
Project policy identifies consequential design and code. No LLM call is spent choosing whether to call an LLM.

3

Independent review
A separate model receives neutral design context or a read-only code snapshot, not the authoring model's persuasive reasoning.

4

Immutable review subject
Base commit, staged and unstaged changes, allowed untracked files, policy version and hashes form one fingerprint.

5

Evidence gate
Material blocks require concrete evidence, consequence and a verification step. Confidence alone is never enough.

6

Protected trust path
Blessed runtime, manifests, owner-bound approval and reviewed maintenance make control weakening visible and fail safely.

COMPETITIVE POSITION

WHAT THE MANAGER SEES

A decision record, not AI theatre

”

Independent review evidence obtained.
Reviewed fingerprint + workspace digest + response hash
Verdict + findings + verification + stale/current status
OWNER: ACCEPT / REJECT WITH EVIDENCE / DECIDE

WHY THIS IS DIFFERENT

- Before and after**
Reviews consequential design before implementation and meaningful code after it.
- Independent by construction**
Neutral design packets reduce anchoring on the authoring model's answer.
- Exact-state assurance**
A review cannot be reused after the code, policy or captured content changes.
- Selective economics**
Routine work remains single-model; duplicate fingerprints do not create repeat calls.
- No automatic repair loop**
The reviewer is read-only. Findings do not silently edit, merge, commit or win.
- Control can be paused**
Late results are quarantined when mode or generation changes.

Most tools review code. 2nd Engineer governs the handoff between AIs.

Vendor-published emphasis, not a claim that unlisted capabilities are absent. Product features change; sources below were checked 1 September 2026.

Product	Primary review point	Trigger model	Published evidence / action	2nd Engineer distinction
2nd Engineer	Pre-design + post-change	Policy-routed GREEN / RED + owner modes	Immutable fingerprint, stale quarantine, structured evidence	Control plane: independent context, exact-state approval, protected trust path
ReviewGate [1]	Agent loop + pre-push + pull request	CLI, hook, MCP or PR automation	Self-hosted policy gate, independent judge and merge status	Closest workflow peer; 2nd Engineer adds neutral pre-design review, owner signatures and protected runtime
GitHub Copilot [2]	IDE + pull request	Manual or auto on PR open, draft or new push	Comments and suggested fixes; review does not count as required approval	Adds an enforceable pre-PR gate and exact-fingerprint owner decision
CodeRabbit [3]	PR + IDE + CLI / pre-commit	Automatic PR review or local command	AI review plus code graph, lint and scanners; suggested fixes	Adds author/reviewer separation, neutral design review and protected control authority
Qodo [4]	Pull request	Git-integrated multi-agent review	Rule enforcement and context-aware findings in the PR	Adds lifecycle control before design and outside the PR, plus stale-state invalidation

Official product sources: ReviewGate product 2. GitHub Copilot code review 3. CodeRabbit review overview4. Qodo code review
Windows + Git | Reversible Claude / OpenAI roles | Team plan: \$179/month for up to 10 developers.
Trust note: threat model excludes a deliberately hostile elevated administrator; residual same-user evidence risk is documented.

TECHNICAL BUYER BRIEF

Architecture, model direction, languages and deployment fit

A concrete view of the reversible model topology, protected control path, language fit and commercial team plan.

MODEL DIRECTION

Author and reviewer are independent, reversible roles.

Claude can author while OpenAI Codex reviews, or Codex / GPT can author while Claude reviews. The model pairing changes; deterministic policy, exact-state evidence and owner authority do not.

AUTHOR ROLE

Claude / Codex / GPT

Selected author produces the design or software change.

CONTROL PLANE

2nd Engineer

Route, capture, fingerprint, attest, gate and quarantine.

REVIEWER ROLE

Codex / Claude

The other model receives the independent read-only contract.

REVERSIBLE MODEL ROLES

Claude -> Codex, or Codex / GPT -> Claude. Policy, evidence and owner gate remain constant.

IMPLEMENTATION AND FIT

The stack is intentionally local, deterministic and inspectable.

RUNTIME AND INTEGRATION

- Node.js 22+ ES modules; no application framework required.
- Git repository provides commits, staged, unstaged and untracked state.
- Role adapters accept Claude Code, Codex and GPT author workflows.
- Protected OpenAI Codex and Claude transports provide independent review.

CRYPTOGRAPHY AND TRUST

- SHA-256 binds content, snapshots, manifests and workspace state.
- Ed25519 verifies protected bootstrap and control artifacts.
- script derives owner-bound protection from a passphrase.
- Windows ACLs protect the ProgramData blessed runtime and trust anchor.

LANGUAGES AND REPOSITORIES

- Capture is language-agnostic because the review subject is Git state plus files.
- Core implementation is built and exercised in Node.js / JavaScript.
- Routing patterns already cover C#, .NET, Blazor, SQL and infrastructure risk.
- Reviewer capability determines semantic depth for each programming language.

DATA AND REVIEW SAFETY

- Reviewer is read-only; it cannot automatically edit, merge, commit or release.
- Snapshot clone has no remote; automatic review uses no network by default.
- Ignored credentials and explicit sensitive paths are excluded.
- Suspected secret material aborts capture; changed results are quarantined.

EXTENSION SURFACES

Portable by adapter, not by wishful thinking

AUTHOR HOOK

request + completion signals

POLICY PACK

paths + terms + control domains

SNAPSHOT

Git capture + secret exclusions

REVIEW TRANSPORT

model-specific invocation + provenance

OUTPUT SCHEMA

verdict + evidence + verification

TEAM PLAN - \$179 / MONTH

ALCHEMYLOOP.AI

Up to 10 active developers. Guided onboarding: \$995 once. Customers supply their own AI provider access; additional teams and integrations are scoped separately.