LEGIT

Agentic Remediation

From Detect to Fixed, Automatically

Legit's remediation agents prioritize, fix, and validate vulnerabilities across every affected service — closing the gap before attackers exploit it.

Overview

AppSec has spent a decade getting brilliant at finding vulnerabilities. The trouble is, finding was never the problem. Fixing issues is. Every finding still lands in the same manual chain: a security engineer triages it, a developer reconstructs the context the scanner threw away, writes a fix, and waits for review. That chain runs at human speed while scanners produce findings around the clock. In this model, the backlog always wins. Legit's agentic remediation closes that gap. Purpose-built agents triage, fix, and validate findings automatically, using context a generic coding assistant doesn't have.

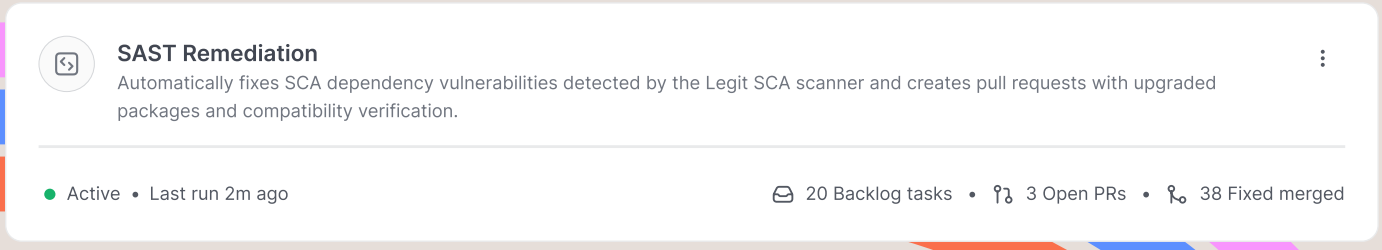

The Problem

- AI-generated code carries 2.74x more vulnerabilities than human-written code.
- The median time to remediate a vulnerability is 252 days, but attackers move from disclosure to exploit in minutes.
- Pointing a general-purpose coding agent at the backlog doesn't fix this. These tools have no security context, can't tell a real risk from noise, and never proves the fix worked.

Agentic Remediation

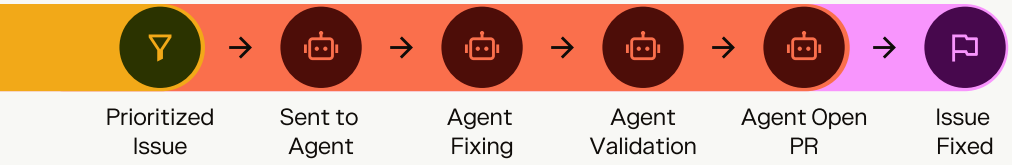

How It Works

Legit's remediation agents run a security-native loop a generic coding assistant can't:

1. **Context.** Legit enriches every finding with reachability, ownership, deployment status, and business impact.
2. **Prioritize.** Reachability and exploitability scoring, org-wide false-positive memory, and business impact strip out the noise so only confirmed risk reaches the queue.
3. **Remediate.** The agent reads, plans, and fixes the root cause — not just the symptom. After that, it opens a pull request routed to the right owner, with a plain-language explanation of what changed and why.
4. **Validate.** The agent re-scans to confirm the finding is closed, with no regressions and no new issues introduced, before anything merges.

What Makes It Different

- Parallel remediation across every affected service. A single vulnerability — like an auth bypass introduced through reused code — often shows up in dozens of services at once. Legit's agents open pull requests across every affected repo in parallel, closing the whole gap instead of one instance of it.
- Root-cause fixes, not patches. Legit's agents carry security expertise and your organization's risk posture — knowledge a general-purpose coding agent simply doesn't have access to.
- Full audit trail. Every action an agent takes — from the original finding to the validated fix to what engineering did with it — is logged for compliance and audit readiness.
- Gets smarter over time. Agents learn from each organization's distinct codebase and every past fix, so remediation gets faster and more accurate with every cycle.

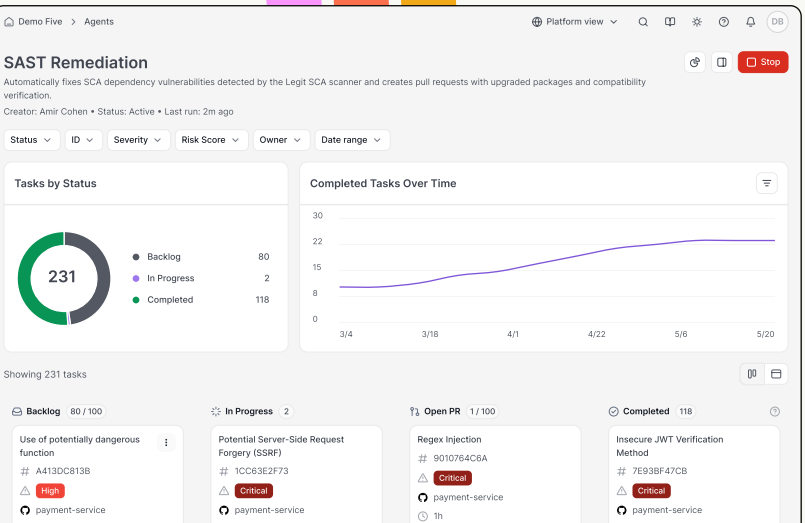

See remediation that doesn't stop at "here's a patch."