

[SOLUTION BRIEF]

Review every PR before it's merged.

Sentinel reviews each pull request the moment it opens and flags exploitable vulnerabilities as inline comments. The signal, not the noise.


[THE PROBLEM]

The bug is rarely in the diff. A linter matches rules against source and reports whatever matches, so it misses the vulnerability that only exists once a change meets the code around it, and floods developers with findings that are not exploitable. Security review that arrives after merge is too late, and review that arrives with noise is ignored.

WHAT SECURITY TEAMS NEED

- 01 Context, not pattern-matching
A reviewer that follows execution paths and untrusted data beyond the diff.
- 02 Findings where developers work
Inline comments on the pull request, each with a fix, before merge.
- 03 Fewer false positives over time
Developer feedback that becomes reusable rules instead of repeated triage.

[BY THE NUMBERS]



BUILT BY THE TEAM THAT SECURED

Apple · Google · Microsoft · U.S. Department of Defense · T-Mobile · Henkel · Epic Games · AT&T · Yahoo

BACKED BY Y COMBINATOR

[HOW IT WORKS]

Connect. Analyze. Review inline.

01

Connect

Connect GitHub, GitLab or Bitbucket Cloud and choose the repositories you want reviewed. Sentinel picks up every new pull request from there.

02

Analyze

Sentinel reads the diff, the surrounding code and the execution paths it touches, following how untrusted data moves through the system before deciding what to flag.

03

Review

Each finding is posted as an inline comment on the pull request with a fix. Developer replies feed back into future reviews and cut false positives.

[SOLUTION ARCHITECTURE]




FIG. 01 Sentinel review loop. A pull request event triggers analysis of the diff plus surrounding execution paths; confirmed, exploitable findings return to the PR as inline comments with fixes, and developer feedback becomes rules.

Understands your code

Sentinel forms a hunch about a change and follows the execution path to confirm it or kill it. If it cannot show you the problem and hand you the patch, it does not speak.


Comments where it matters

Findings land directly on the pull request as inline comments. Create a Jira or Linear ticket from a comment when work needs to be tracked.

Improves with every review

When a developer marks a finding wrong, that feedback is captured as a reusable rule that sharpens later reviews across the organization.

Exploitable classes, caught at the diff.



[TECHNICAL SPECIFICATIONS]

SOURCE CONTROL	GitHub, GitLab, Bitbucket Cloud
TRIGGER	Every new pull request on selected repositories
LANGUAGES	JavaScript, TypeScript, Python, Go, Java, Kotlin, Rust, Ruby, PHP
ANALYSIS SCOPE	Diff plus surrounding code and the execution paths it touches
OUTPUT	Inline PR comments with explanation, confidence and fix
TICKETING	Create Jira or Linear tickets from a comment
FEEDBACK LOOP	Developer replies captured as reusable Rules
DEPLOYMENT	SaaS integration with your Git provider; no agent to install

BUILT FOR MONOREPOS

- **Diff-scoped analysis**
Sentinel reviews the diff and the execution paths it touches, not the entire codebase on every pull request, so review times stay fast in large monorepos.
- **Nine languages, common frameworks**
JavaScript, TypeScript, Python, Go, Java, Kotlin, Rust, Ruby and PHP, with the common frameworks and libraries in each.
- **Fix included**
Every finding ships with a patch. Developers act on the comment without leaving the PR.
- **Read-only by design**
Sentinel needs read access to review and comment permissions to post. It never pushes to your branches.

[VULNERABILITY CLASSES]

AUTH Authentication bypass	AC Broken access control	IDOR Insecure direct object reference	LOGIC Business logic flaws
INJ Injection	SSRF Server-side request forgery	SECRET Leaked secrets	INPUT Unsafe handling of user input

[WHY PARAMETER]

- ✔ **Exploitable, not pattern-matched**
Sentinel flags issues that are actually exploitable rather than matching on keywords, and proves it before it comments.
- ✔ **312,940 PRs in 30 days**
9,417 vulnerabilities caught in the last 30 days across customer repositories, about one review every 17 seconds.
- ✔ **Every finding ships with a fix**
If Sentinel cannot show the problem and hand you the patch, it does not speak.
- ✔ **Same team, same platform**
Findings correlate with the running-app pentest, cloud and dependency modules in one queue.

[PART OF THE PARAMETER PLATFORM]

Sentinel closes the loop with the Pentesting Agent: a class of bug proven in the running application becomes a pattern Sentinel watches for at the diff, before the next one merges.

See first findings in 24 hours.

ABOUT PARAMETER

Parameter is an AI-powered security platform that continuously finds and validates vulnerabilities before attackers can exploit them. Autonomous agents pentest like a real adversary across your code, cloud and dependencies, with every finding proven. Formerly Hex Security, Parameter is backed by Y Combinator and built by a team that has secured Apple, Google, Microsoft, the U.S. Department of Defense, T-Mobile, Henkel, Epic Games, AT&T and Yahoo.

TRUST AND COMPLIANCE

SOC 2 Type I (Type II observation underway)
HIPAA BAA available
GDPR DPA with SCCs
Data processed in the US
parameter.ai · support@parameter.ai · app.parameter.ai