

AI SECURITY ENGINEERING

Security that keeps up with engineering.

AI security agents that audit your codebase, prove what is exploitable, and hand your engineers fixes they can review and merge.



BACKED BY Y COMBINATOR

SOC 2 TYPE II

SSO · ZERO DATA RETENTION

WINFUNC.COM

Your code is growing faster than anyone can review it.

90%

of enterprise software engineers will use AI code assistants by 2028, up from under 14% in early 2024.

Gartner, 2025

45%

of AI code-generation tasks introduced an OWASP Top 10 vulnerability into the code.

Veracode GenAI Code Security Report, 2025

31%

of breaches now begin with vulnerability exploitation, the top entry point for the first time in 19 years.

Verizon DBIR, 2026

More code, written faster, reviewed by the same number of security engineers.

Every option you have today leaves a gap.

PATTERN-BASED SCANNERS

Fast and broad.

Good at known patterns across every repository, on every build.

THE GAP

Most alerts are noise. In a study of 2,856 open-source repositories, **91.5% of findings were false positives.**

Ghost Security, 2025 (vendor study)

PENETRATION TESTS

Deep and creative.

Human judgment finds attack paths that rules never will.

THE GAP

A snapshot. The report describes code that has already changed by the time it arrives, and the next test is a year away.

INTERNAL SECURITY REVIEW

Knows your product.

Your own engineers understand the business logic best.

THE GAP

It doesn't scale with commit volume, and security specialists are the hardest hires to make.

TRIAGE LOAD, WORKED EXAMPLE

30 findings a week × 20 minutes each × \$100 an engineer-hour

43

hours a month, about \$52K a year, spent confirming alerts

An AI security engineer that shows its work.

Winfunc's agents work through your codebase the way a senior security engineer would, and leave a record of every step.

01 Understand Map entry points, roles, tenants, trust boundaries, and data flows. Record Reviewed scope and revision	02 Hypothesize Generate attack ideas specific to your application, including business logic. Record Candidate attack paths	03 Investigate Follow each path from attacker-controlled input to the vulnerable operation. Record Source-to-sink trace	04 Prove Check reachability and existing controls; reproduce where it is feasible and in scope. Record Validation status	05 Fix Write a patch in your codebase's style and open it as a pull request. Record Suggested patch PR	06 Verify Re-review the updated code and return a clear fixed or still-present verdict. Record Re-review verdict
--	--	---	--	--	--

YOUR CONTROL

Engineers approve every change. Winfunc never merges to a protected branch, and untested impact is labeled as inference.

One workspace, from finding to merged fix.

Security score per repository

0 to 100, weighted by severity, how long issues stay open, and scan cadence.

Findings, aging, and triage

Open findings by severity, how long each has waited, and where it stands in triage.

Autofix per repository

See where autofix is on, and which repositories have findings waiting for a patch.

Coverage and burndown

Functions analyzed, coverage score, and open findings over time.

winfunc

Dashboard

Repositories

Settings

Updates

22

AS

acme-security

security@acme.dev

OVERVIEW

Security posture

Last scan 21m ago

+ ADD REPOSITORIES

OPEN FINDINGS

44

-12 vs previous

CRITICAL + HIGH

14

oldest 9d

C 3

H 11

SECURITY SCORE

82

+8 vs previous

82

REPOSITORIES

5

all fresh

132 scans

AWAITING TRIAGE

9

12 resolved

P 6

V 3

2 repositories have findings but autofix off · acme/web-platform, acme/infra

NEEDS ATTENTION 4 of 5 repositories

ALL REPOSITORIES

Repository	Score ↑	TREND	Findings	Last scan	Autofix
acme/commerce-api	62 +9	↗	C 2 H 5 M 8 L 3	21m ago	ON
acme/payment-service	74 +5	↗	C 1 H 3 M 5 L 2	37m ago	ON
acme/web-platform	86 +1	—	C 0 H 2 M 4 L 2	1h ago	OFF
acme/identity-service	91 +4	↗	C 0 H 1 M 2 L 1	2h ago	ON
acme/infra	96	—	C 0 H 0 M 1 L 2	3h ago	OFF

5 REPOSITORIES

H

FINDINGS 44 open

C 3

H 11

M 20

L 10

AGING 44 unresolved

< 7 DAYS

7-30 DAYS

30-90 DAYS

> 90 DAYS

14

17

9

4

TRiage 9 awaiting

PENDING

VALIDATING

ACCEPTED

WON'T FIX

RESOLVED

6

3

35

2

12

COVERAGE 91.5%

FUNCTIONS ANALYSED 4,412 / 4,820

PRODUCT UI · SAMPLE DATA

05 / 19

See exactly how each finding was reached.

CRITICAL 9.8 RAY · RAY-PROJECT/RAY FIXED AUG 2026

Unauthenticated Serve replica gRPC requests allow arbitrary code execution

SOURCE src/ray/protobuf/serve.proto:525 · ASGIRequest.pickled_request_metadata
SINK python/ray/serve/_private/replica.py:248 · _wrap_grpc_call.wrapper()
CVSS 3.1 / AV:N / AC:L / PR:N / UI:N / S:U / C:H / I:H / A:H

1 Replica.__init__()
Each replica creates a raw gRPC server outside Ray's authentication factory.

```
self._server = grpc.aio.server(options=grpc_options)
```

2 _on_initialized()
The service listens on every interface with no TLS or token check.

```
port = self._server.add_insecure_port("[::]:0")
```

3 _wrap_grpc_call.wrapper()
Attacker-supplied bytes are deserialized before any authentication.

```
request_metadata = pickle.loads(request.pickled_request_metadata)
```

VALIDATION

Reproduced on an isolated cluster with token auth enabled. A request carrying no token reached the pickle sink.

OUTCOME

Upstream fix merged 5 Aug 2026 (PR #65189): replicas now use Ray's authenticated, TLS-aware gRPC helpers.

WHAT EVERY FINDING CARRIES

SCOPE The repository, revision, and authorized boundaries that were reviewed.

TRACE The path from attacker-controlled input to the vulnerable operation, line by line.

VALIDATION What was reproduced, which controls were checked, and what remains inference.

SEVERITY A CVSS vector with the reasoning behind it.

FIX A suggested patch as a pull request, then a re-review once the code changes.

Proof your engineers can rerun.

Every finding ships with a proof of concept your engineers can run themselves, plus a validation run that states its confidence.

ENVIRONMENT	How to stand up the vulnerable target locally.
STEPS	The exact requests or commands that trigger the issue.
EXPECTED	What you will see when the exploit works.
VALIDATION	An independent run records confidence and any remaining risk.

>_ PROOF OF CONCEPT

COPY POC

ENVIRONMENT

BASH

1 docker compose up -d api postgres

2 bun run seed:demo # creates organizations 1041 and 1042

STEPS

BASH

1 # 1. Log in as a member of organization 1042

2 curl -s -c jar.txt -X POST localhost:4000/api/auth/login \

3 -d '{"email":"member@org-1042.test","password":"demo-password"}'

4

5 # 2. Export another tenant's orders

6 curl -s -b jar.txt "localhost:4000/api/orders/export?organizationId=1041&format=csv" |

EXPECTED RESULT

TEXT

1 order_id,customer_email,shipping_address,total


2 A-10231,jane@org-1041.test,"12 Harbor St, Oakland",184.20

3 A-10232,omar@org-1041.test,"88 Pine Ave, Denver",62.00

Fewer findings, and they hold up.

Winfunc reports what it can support with evidence. Its research results are public, including the misses.

PATTERN-BASED SAST, AT SCALE



180 8.5%

real vulnerabilities

1,936 91.5%

false positives someone had to triage


Worst case in the same study: command injection in Python/Flask, where 1,160 of 1,166 findings were false positives.



6 real · 99.5% false positives


Ghost Security, "Exorcising the SAST Demons," 2025 · 2,116 findings across 2,856 open-source repositories · vendor study

WINFUNC, PUBLISHED RESULTS



6/6

NGINX findings fixed upstream, with five CVE IDs credited by F5. Jul 2026.



10/12

held up under live reproduction in the HackerNews codebase audit. The two misses are published too. Apr 2026.

13

bugs patched upstream across nine projects, including Node.js, React, and NGINX, with seven CVEs. Mar 2026.

Different datasets, shown side by side for scale. This is not a head-to-head benchmark. In your workspace, every finding carries its own validation status.

Findings arrive with the fix.

- 01

Patch generated from the trace

Written in your codebase's style to close the path that was proven.
- 02

Opened as a pull request

Autofix tracks the branch and PR next to the finding.
- 03

Validation checks the patch

An independent run marks it Patch valid, with confidence and remaining risks.
- 04

Your engineers review and merge

Nothing merges without your approval. PR re-reviews return a clear verdict.

● Fixed

● Still present

📄 PATCH 1 file +5

UNIFIED SPLIT

📄

📄

📄 src/routes/orders.ts +5

↕

ROUTER.GET("/EXPORT", REQUIREAUTH, ASYNC (REQ, RES) => {

58

const organizationId = req.session.orgId;

59

const organizationIdrequested = Number(req.query.organizationId ?? organizationId);

60

61

if (requested !== organizationId && !can(req.user, "orders:export:any")) {

62

return res.status(403).json({ error: "forbidden" });

63

}

64

const format = parseFormat(req.query.format);

65

const stream = await exportOrders({ organizationId: requested, format });

66

res.setHeader("Content-Type", contentType(format));

67

stream.pipe(res);

68

});

🛡️ VALIDATION Latest run

● PATCH VALID

CONF 97.0%

a1b2c3d

1h ago

REMAINING RISKS

▪ Admin export endpoint shares the same service; review after the fix lands.

ANALYSIS

Reproduced against the seeded demo environment. A member of organization 1042 exported organization 1041's orders; the response contained 212 rows with customer emails and shipping addresses. No session-scoping check exists on the export path.

○ PRODUCT UI · SAMPLE DATA

09 / 19

Shorten time-to-fix by removing the handoffs.

INDUSTRY AVERAGE

252

DAYS TO FIX A FLAW

Up 47% since 2020. Leading teams fix half their flaws within five weeks; lagging teams take more than a year.

Veracode, State of Software Security 2025

WHERE THE TIME GOES

A tool flags a possible issue

SCANNER

Someone confirms it is real

WINFUNC

Someone reproduces it and explains the impact

WINFUNC

Someone writes the first draft of a fix


WINFUNC

An engineer reviews and merges

YOUR TEAM

Winfunc delivers the middle three together, so your team starts at review.

MEASURED IN YOUR WORKSPACE



Oldest open critical or high

The number leadership asks about first.

Fixed before merge

Share of PR findings closed before they ship.

We baseline these in the first week of a pilot and report the change.

Ask the security question that keeps you up at night.

- ? Can a user access another tenant's exports?
- ? Can webhook signatures be bypassed?
- ? Can an invite be replayed after acceptance?
- ? Can a read-only role trigger a write?

TRACED

The code paths, roles, and state behind the guarantee.

WEIGHED

Evidence for and against, tied to the revision.

BOUNDED

Limits stated plainly. No finding is not proof of safety.

ACTIVE SCAN

15m 02s

RUNNING

DETAILS

STOP

Can a read-only role trigger a write through the bulk-update endpoint?

BRANCHmain

OVERALL

live

ACTIVITY

waiting for stream

Running focused attack missions...

MAX CHARGE\$6.5

RESERVED\$6.5

CHARGED\$0

FINDINGS0

RECENT SCANS

4 on record

STATUS	HYPOTHESIS	CREATED	DURATION	FINDINGS	CHARGED	
COMPLETED	Can a user access another tenant's exports?	46m ago	19m 00s	2	\$7.9	FINDINGS ↗
COMPLETED	Can webhook signatures be bypassed?	2d ago	14m 00s	0	\$5.1	FINDINGS ↗
COMPLETED	Can an invite be replayed after acceptance?	4d ago	22m 00s	1	\$6.3	FINDINGS ↗
COMPLETED	Can a refund exceed the captured charge?	6d ago	17m 00s	1	\$5.8	FINDINGS ↗

4 SCANS

Everything in the platform today.

AUDIT Security audit A version-bound review of your codebase with traced, validated findings. For a launch, a raise, or an enterprise deal	CONTINUOUS Pull-request security Every change reviewed against your repository's rules, with fixes re-verified. Plus scheduled scans	INVESTIGATE Hypothesis scans Deep investigation of one security question, asked in plain words. Business logic and authorization	SUPPLY CHAIN Dependency scanning OSV-backed results with fixed versions, upgrade commands, and rescans. npm, PyPI, Go, Maven, and more
REMEDiate Autofix pull requests Patches written in your codebase's style and opened for review. Re-review verdict on every fix	TRIAGE Triager assistant Ask anything about a finding; the assistant knows your repository. Validation and revalidation	REPORT Analytics and reports Security score, aging, burndown, and printable reports for audits. Per organization and repository	IN THE EDITOR Winfunc MCP Check a snippet, file, or before-and-after change from your AI editor. Cursor, Claude Desktop, Windsurf, Cline

AVAILABLE	Threat modeling · Cloud security	COMING NEXT	Triage for other tools' findings · Leadership views
-----------	----------------------------------	-------------	---

Found in software you already run.

48

PUBLIC DISCLOSURES

22

CVE IDS

25

PROJECTS

33

CRITICAL OR HIGH

RESEARCH · JUL 2026

Six NGINX vulnerabilities found with open models. Five CVE IDs, credited in F5's advisories.

BENCHMARK · MONTHLY

N-Day-Bench tests how well frontier models find real vulnerabilities disclosed after their training cutoff.

PROJECT	FINDING	SEVERITY	REFERENCE
NGINX	Stream capture desynchronization causes a heap overflow	HIGH	CVE-2026-42533
Chromium	Type confusion in V8 (details withheld)	HIGH	CVE-2026-10910
Ray	Unauthenticated Serve replica gRPC allows code execution	CRITICAL	PR #65189
Anthropic FastMCP	Authentication bypass on custom routes	CRITICAL	Disclosed
Supabase	SQL injection through queue names	CRITICAL	Disclosed
Gumroad	Zero-click account takeover through a helper endpoint	CRITICAL	Disclosed
React	Server Components reply decoder denial of service	HIGH	CVE-2026-23864
Node.js	Permission model bypass through Unix domain sockets	MEDIUM	CVE-2026-21636
Mattermost	Ten issues, including a role-granting authorization bypass	HIGH	10 CVEs

How Winfunc compares.

	Winfunc AI security engineering	Pattern-based SAST e.g. Semgrep, Snyk Code, CodeQL, Checkmarx	AI code review bots e.g. CodeRabbit, Copilot code review	Runtime AI pentest e.g. XBOW, RunSybil	Human pentest Security consultancies
WORKS FROM	Your code plus application context	Code patterns and rules	The pull-request diff	The running application	An agreed scope and time box
BUSINESS LOGIC	Investigated through roles, tenants, and state	Only where someone wrote a rule	Limited to what the diff shows	What is observable from outside	Strong, within the engagement window
EVIDENCE	Source-to-sink trace, validation status, reproduction where feasible	A rule match and a line number	A review comment	Request and response proof	A written report
FIX	Patch PR, then re-review	Guidance, some autofix	Suggested change	Guidance	Recommendations
CADENCE	Every PR, on a schedule, or on demand	Every build	Every PR	Scheduled or continuous	Yearly or per release
YOUR TEAM TRIAGES	Evidence-backed findings	High alert volume	Review comments	Runtime-confirmed issues	One report per engagement

KEEP WHAT WORKS

Winfunc complements a pentest program: keep the pentest for human judgment at a point in time, and use Winfunc for code-level coverage as the code changes. Ask us for a side-by-side on your own code.

What engineering leaders say.



There were vulnerabilities that we thought we took care of that the Winfunc agent found extremely complex bypasses for, vulnerabilities that other security tools never would have found.



Noah
Co-Founder & CEO, Scout



Dennis
Co-Founder & CEO, Surge



Ram
Co-Founder, Sei



Bereket Engida
Founder, Better Auth



Enterprise-ready from day one.

Single sign-on, zero data retention, customer-controlled deployment, and a complete audit trail, backed by a SOC 2 Type II audit.

IDENTITY AND ACCESS	DATA PROTECTION	DEPLOYMENT AND MODELS	OVERSIGHT
✓ SAML 2.0 single sign-on Okta, Microsoft Entra ID, Google Workspace, or any SAML 2.0 identity provider.	✓ Zero data retention Source code is analyzed in ephemeral workspaces and deleted after each scan. Only findings are kept.	✓ Customer-controlled hosting ENTERPRISE Run Winfunc inside your own cloud or private environment.	✓ Complete audit trail Every API request and authorization decision is logged for administrative review.
✓ Federated login and provisioning Manage who gets access from the identity provider you already run.	✓ No training on your code Customer code and data are never used to train AI models.	✓ Isolated tenants ENTERPRISE Dedicated single-tenant environments for regulated workloads.	✓ Least-privilege access Winfunc reads only the repositories you select and never merges to protected branches.
✓ Repository-scoped roles Owners decide who sees findings and who can run scans on each repository.	✓ Sealed credentials TLS on every connection. Test credentials are encrypted with per-record keys and rotatable master keys.	✓ Bring your own model keys ENTERPRISE Route analysis through your own AI provider accounts and agreements.	✓ Contract protections DPA and HIPAA BAA available. Security incident notice within 72 hours.
✓ Revocable access keys API and MCP keys are stored only as SHA-256 hashes and can be revoked at any time.	✓ Secret-free logs Logs keep request metadata only: never bodies, tokens, cookies, or authorization headers.	✓ Open-weight model support Proven in public research: the NGINX work that produced five CVE IDs ran on open GLM models.	✓ Named support ENTERPRISE A named security contact with contracted response terms.

Works where your engineers already work.

SOURCE CONTROL



GitHub
App · repos, PRs, patches



GitLab
repos and merge requests



Bitbucket
repository workflows



Azure DevOps
Azure Repos and Pipelines

CI/CD



GitHub Actions
winfunc-action



Jenkins
pipeline scans and thresholds



CircleCI
workflow scans

TICKETS AND ALERTS



Jira
issues from reviewed findings



Slack
scan and finding alerts



Microsoft Teams
notifications and summaries

AI EDITORS · WINFUNC MCP



Cursor
via Winfunc MCP



Claude Desktop
via Winfunc MCP



Windsurf
via Winfunc MCP



Cline
VS Code, via Winfunc MCP

LANGUAGE COVERAGE

Parsers, language servers, and model-assisted analysis, so coverage isn't tied to one language's rule set. Public findings already span:

- C
- C++
- Go
- Rust
- Python
- TypeScript
- PHP
- Ruby
- Lua

SETUP

Connect a repository through the app for your source-control platform, choose the branches and pull-request workflows to cover, and set repository rules for what each scan should focus on and report.

Start with one audit. Keep what proves itself.

WAYS TO BUY

\$500+	One-time audit A surface scan from \$500, or a deep scan priced by scope. No subscription.
Credits	Pay as you go Buy credits when you need them and use the full workflow.
Monthly	Recurring coverage A renewing credit allowance for continuous PR review; discounted annually.
Custom	Enterprise SSO, customer-controlled hosting, your own model keys, and named support.

Credits are estimated before each scan from changed files, lines of code, and comparable scans, then reconciled to actual usage.

YOUR FIRST TWO WEEKS

DAY 0

Connect repositories and agree on scope

Pick the application, branches, and the questions that matter most.

DAY 1

First findings, with evidence

Sei had a detailed scan in less than a day. Timing depends on repository size and depth.

WEEK 1

Readout with your engineers

Walk through each finding; patches arrive as pull requests.

WEEK 2

Turn on continuous coverage

PR review and scheduled scans, with time-to-fix metrics baselined.

Bring us your hardest security question.

We'll run a hypothesis scan on it against your code and walk your engineers through the evidence.

WHAT YOU GET

A traced answer to your question

Findings with proof and suggested fixes

A baseline for your time-to-fix metrics

[Book a demo at winfunc.com →](#)

CONTACT

EMAIL
hello@winfunc.com

WEBSITE
winfunc.com

TRUST PORTAL
winfunc.trust.site

PUBLIC FINDINGS
winfunc.com/hacktivity

