ShipShape

Find the security holes in your app before a buyer does.

A security scan built for people who shipped something real with AI tools, and never had anyone check it. Plain English, no jargon, no security background needed.

THE PROBLEM

Anyone can build software now. Not everyone can secure it.

People are shipping real products in a weekend using AI tools. Those products hold customer data, API keys and payments, and most of them have never had a security review of any kind. Not out of carelessness. There just hasn't been an option built for a non-specialist.

WHO IT HELPS

One person, not "developers" broadly.

The solo builder

Built something real with AI tools. Has actual users. No security team, no budget for a professional penetration test, and no way to read the output of the tools built for experts.

The moment they need it

A customer sends a security questionnaire. A launch is coming up. Or they read about a breach and get the quiet feeling they might have left a door open somewhere.

Later: the agencies and freelancers shipping client apps, who want a security check they can hand over as part of the work.

WHAT IT DOES

Paste a URL. Get a report you can actually read.

- ✓ **Exposed secrets** — API keys and credentials sitting somewhere the public can reach them
- ✓ **Configuration gaps** — missing security headers, insecure cookies, unencrypted connections
- ✓ **Cross-site scripting** — whether a stranger can run their own code in your users' browsers
- ✓ **Broken access control** — whether changing a number in a URL shows someone else's data
- ✓ **Prompt injection and AI data leaks** — whether your AI features can be talked into misbehaving

Every finding comes with what it is, why it matters, and one concrete thing to change.

WHAT MAKES IT DIFFERENT

It tests the AI holes the other scanners miss.

1

AI-era vulnerabilities

Prompt injection and AI data disclosure. Most security tools were designed before anyone shipped an app with a language model inside it, and they simply do not look for this.

2

Written for the builder

The report is written for the person who made the app, not for a security engineer. No severity codes to decode, no raw request dumps to interpret.

THE PRINCIPLE

It never claims more than it checked.

A security tool that overstates what it looked at is worse than no tool at all, because people act on it. Every report lists exactly which classes were tested and which were not. A partial scan is never presented as a clean bill of health, and the tool will not run its deeper tests against a site until you have shown it is yours.

That restraint costs conversions. It is the whole product.

HOW PEOPLE FIND IT

Meet builders where they already are.

Social and paid

Short video showing a real scan finding real problems, promoted across social platforms and paid social including Meta.

Creator partnerships

The people teaching others to build with AI tools have exactly the right audience, and a recommendation from them carries far more weight than an ad.

The report itself

A shareable result that gets forwarded to a customer or a co-founder is the cheapest and most credible marketing there is.

Start small, measure honestly, and put the money behind whatever actually works.

WHERE IT IS TODAY

Working, live, and honest about being early.

Built and running

Both scan modes work end to end. Accuracy is graded automatically on every change, so a regression that starts producing false alarms fails the build rather than reaching a user.

Next up

Broader coverage for single-page apps, a shareable report link, and turning findings into something a customer's security questionnaire can be answered with.

Security shouldn't require a security background.

ShipShape is for everyone building real things with AI tools who wants to know, honestly, whether they left a door open.

Not a certification. It reports only the classes it tested, and an absence of findings is not proof that an app is safe.