

PRODUCTION HARDENING SPRINT · SEED-STAGE SOFTWARE COMPANY · INVESTOR TECHNICAL REVIEW

Ready for an investor's technical review in ten working days

ABOUT THIS DOCUMENT

This page shows what we (AxonBuild) did for Crewproof, Inc. in a 10-working-day Production Hardening Sprint. It is for founders and for the investors who review their apps. First read what was wrong, then the table of what is true now. Crewproof's due diligence pack holds the evidence for every result.

Crewproof, Inc. makes an operations app for commercial cleaning contractors. Managers plan the jobs, crews record the work on their phones, and the contractors' own customers sign off the work. Maya Okafor (CEO and co-founder, not technical) built the first version in Lovable, an AI app builder. Luis Ferreira (part-time contract developer) changed it later with Cursor, an AI coding tool. The app runs on Supabase (database), Stripe (payments) and OpenAI (AI). 41 contractor companies pay for it, and about 1,500 people use it. Crewproof had no CTO. Its lead investor asked for a technical review (technical due diligence) before the seed round could close. The investor's technical advisor sent about 40 technical questions.

What we found

The app worked well for customers, but nothing kept each contractor's data apart. These were the 6 critical problems:

- A user of one contractor company could read the invoices of every other company.
- Any user could make themselves an admin with one hidden request to the server.
- Any signed-in user could use the admin actions.
- The OpenAI key (the password for the AI service) was in the code that every browser downloads.
- Anyone with the link could open the job photos, which show client buildings and crew faces.
- A fake payment message could mark an invoice as paid.

We checked the app against all 123 items of the sprint. We found **6 critical, 30 high, 51 medium and 26 low** problems. Uploaded files had no backup. Every change went straight to the live app. There were no automatic tests, and nobody was told when errors happened.

What changed

This table compares the app before and after the sprint. Read the green column first: it shows what is true now.

TOPIC	BEFORE THE SPRINT	AFTER THE SPRINT
Data separation	On 3 tables, one company could read another's data.	All 24 tables block it. 312 tests check it on every change.
Exposed keys	5 keys (system passwords) sat in the code or its history.	We replaced all 5 keys. 0 keys are exposed now.
Disaster recovery	Nobody had ever tested a restore of the data.	A drill brought the app back in 48 minutes.
Speed under load	At 150 people at once, 95 of 100 requests took up to 3.9 s.	At 400 people at once, 95 of 100 finish within 0.41 s (410 ms).
Test copies	Only the live app existed (production only).	Local, staging (test) and production (live) copies exist.
Automatic tests	The app had 0 tests.	412 tests run on every code change (CI).
Service accounts	Crewproof owned 5 of its 9 accounts with outside services.	Crewproof owns all 13 of 13. Each needs 2FA (a second sign-in).

How the ten working days went

Days 1 to 2	Days 3 to 7	Days 4 to 6	Days 6 to 8	Day 8	Days 9 to 10	3 working days later
We checked all 123 items. We fixed the critical problems first.	We fixed and tested all week. We added test copies, automatic checks and alerts.	In the same days, a load test (many users at once), an attack test, an upload test and a restore drill all failed at first.	We fixed each cause, and all four retests passed. The report keeps both results.	The second recovery drill took 48 minutes.	We wrote the documents. We handed over live, with the investor's advisor present.	At the review call, every question had a written answer.

What Crewproof received

- An 81-page due diligence pack with the readiness report, architecture, data model, controls, capacity, vendors and licences.
- Rules for AI coding tools and 7 automatic checks. The team can keep building with AI safely.
- 10 operating runbooks (step-by-step guides for incidents and routine jobs). We walked through or rehearsed each one.
- A recorded handover and code tour. We fix defects for 14 calendar days and answer questions for 30 calendar days.

123 of 123 sprint items are closed. 122 have evidence. 1 does not apply, and the report gives the reason. 0 are open.