ONE-PAGE SUMMARY

Security Review of a Live AI Portal

A field-operations portal on Next.js and Supabase, with AI drafting, document extraction and an in-app assistant. Reviewed as a hostile user and a hostile email would see it.

SEPTEMBER 2026 REVIEW · NOVEMBER 2026 FOLLOW-UP · CLIENT UNNAMED

ISSUES FOUND	HIGH SEVERITY	FIXED AND LOCKED BY CI	KNOWN-VULNERABLE PROD PACKAGES
9 + 1 proof	4	All	0

I WHAT WAS TESTED

AI features, inbound email, webhooks, who can see and change what, and third-party packages. Every issue was reproduced in an isolated copy of the app, fixed in code, then covered by an automated test that fails the build if it returns. No test touched production data.

II FINDINGS

A2	The AI document reader could be pointed at outside servers (SSRF).	High	Fixed
B1	Any signed-in staff member could send email from the company address.	High	Fixed
B2	Incoming webhooks accepted forged or replayed requests.	High	Fixed
M10	Nothing checked for vulnerable packages or committed secrets.	High	Fixed
A1	A crafted inbound email could steer an AI-drafted reply.	Medium	Fixed
A3	The AI assistant trusted the browser about the user's role.	Medium	Fixed
A4	AI features had no usage limit, so the API bill could be run up.	Medium	Fixed
R1	November: a later database change re-opened draft-report writes to non-admin staff.	Medium	Fixed on production
B4	The browser security policy was missing four standard controls.	Low	Fixed
B3	Proof check: staff cannot raise their own permissions.	Info	Verified

III WHAT KEEPS THEM FIXED

- A regression test per finding, run on every change.
- 5 database permission suites, plus a canary that must fail.
- 63 end-to-end tests on a local copy of the database.
- 14 AI routes under 15 guard tests: no database access, sign-in required.
- A package and secret scan that blocks high and critical issues.

IV STILL OPEN, TRACKED

- Prompt-injection defenses lower the risk. They can't remove it.
- One AI summary still needs the same untrusted-input handling.
- Attachment downloads in inbound email are a hardening target.
- The browser policy still allows inline scripts.
- 11 flaky end-to-end tests are parked, each with a written reason.

VERDICT

Ready for sensitive records. The open items above are tracked in the full report.