

PAUL WARNER

AI IMPLEMENTATION ARCHITECT

pwarn333@gmail.com | [GitHub](#) | [Portfolio](#)

PROFESSIONAL SUMMARY

Evidence-backed portfolio centered on multi-surface application architecture, secure data boundaries, governed agent infrastructure, and explicit human authority over consequential actions.

CORE CAPABILITIES

Application and workflow architecture | Authentication and authorization | Row-level security | Fail-closed integration boundaries | Production-build and browser validation

SELECTED SYSTEMS

KEYTURN CLEANING COMPANY OS

Production Prelaunch — Controlled Launch | Private canonical repository; sanitized evidence

- Multi-surface operating system connecting a public cleaning website, protected CRM, crew portal, Supabase data boundaries, and guarded automation controls.
- Repository evidence supports typed agent tools, budgets, audit records, explicit approvals, containment, and deterministic behavior while providers are disabled.
- Current limit: production prelaunch does not mean live operations; messaging, payments, and autonomous external actions remain disabled pending owner approval.

GOVERNMENT CONTRACT INTELLIGENCE PLATFORM

Functional Prototype / In Development — Not Deployed | Private canonical repository; sanitized evidence

- Private multi-tenant prototype that preserves federal source records, normalizes opportunities, and applies explainable qualification rules.
- Deterministic pursuit briefs and source-grounded reporting keep GO / HOLD / NO-GO decisions under human control; optional model summaries remain bounded.
- Current limit: no production service, billing, email provider, production database, or autonomous submission workflow is represented as active.

AGENCY OPERATIONS OS

Functional Prototype / In Development — Local Only; Not Deployed | Private canonical repository; sanitized evidence

- Local prototype covering lead intake, qualification, client and project records, approvals, change requests, support, and organization-scoped portal access.
- Repository evidence supports authentication, server-side role checks, organization membership boundaries, row-level security, and browser/database validation.
- Current limit: no AI runtime or production deployment is verified; the system is represented as operations and integration evidence.