

Technical Architecture & Implementation Plan

Architecture, technology stack, tool implementation details, security design, infrastructure, and delivery specifications for the RealizeAI Value Micro-SaaS Platform. Updated to reflect Developer Brief V2.0 (March 2026) with V4 hosting-first pricing. Delivered under a Build + Operate engagement.

Platform: realizeaivalue.com

Stack: React · Supabase · Stripe · OpenAI · Puppeteer

MVP Tools: 17 (Categories 1-2)

Duration: 5 Days (1 Week)

Date: 24 March 2026

INTERACTIVE COMPANION



Volume II: Technical Architecture

Architecture overview, pipeline visualization, 20-tool matrix, database schema, security controls, and delivery map.

<https://proposal-viewer-42380604425.us-central1.run.app/p/a89dc9963e81194a/vol2>

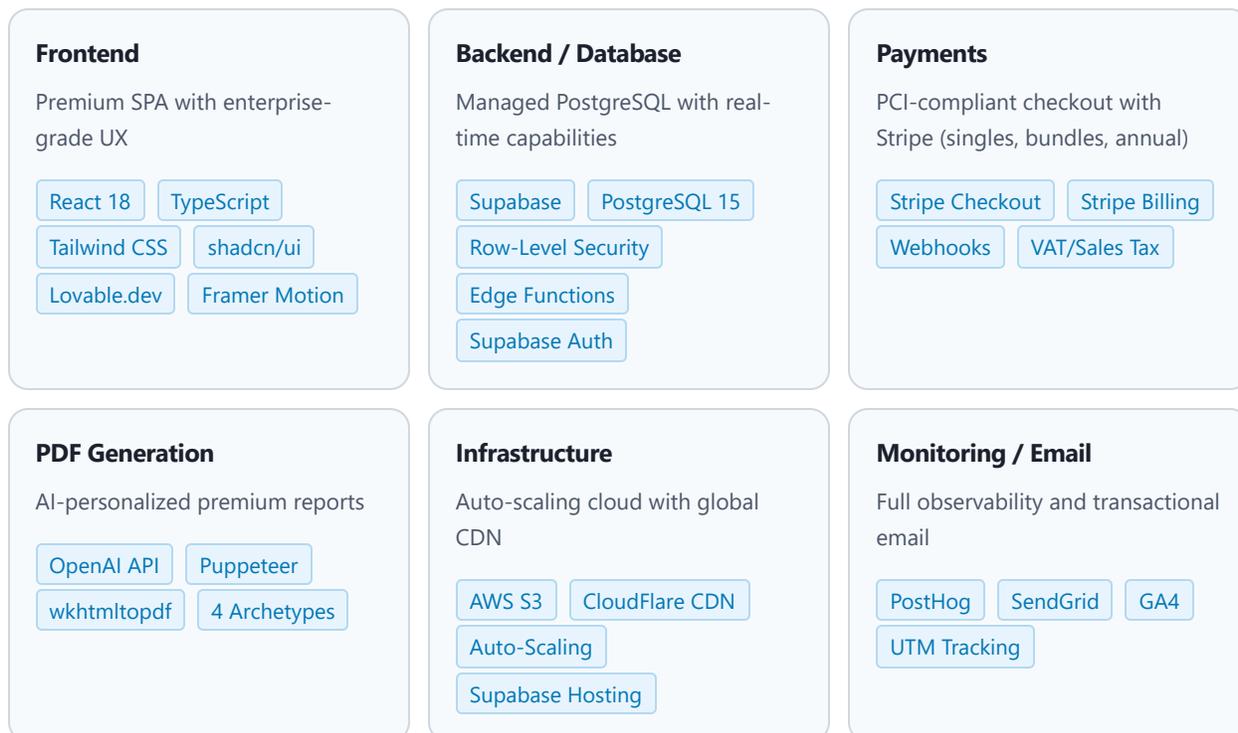
Architecture Overview

The RealizeAI Value platform follows a **modern JAMstack architecture** with a React SPA frontend, Supabase backend-as-a-service (PostgreSQL + Auth + Storage), Stripe for payments, and an OpenAI-powered PDF generation pipeline. The V2.0 architecture prioritizes a **Form-Before-Download** user flow with mandatory data capture for benchmarking, personalized output via maturity archetypes, and a **6-Category Hub** navigation model (MVP: Categories 1-2, with Categories 3-6 visible as "Coming Soon").

Delivery Model: All architecture and integration work is led by **John Lewis** (Solutions Architect), augmented by BFI's 13-agent AI fleet. Agents handle code scaffolding, design system generation, QA automation, CI/CD, and documentation — enabling a 5-day delivery timeline for 20 tools that would traditionally require 24+ weeks with a 5-person team. See Volume I for the full delivery model and Volume III for detailed BOE.

V2.0 Architecture Changes: This volume reflects the Developer Brief V2.0 updates including: (1) 6-Category Hub navigation replacing 5-Stage Pipeline, (2) Form-Before-Download access model replacing credit-based gating, (3) OpenAI API + puppeteer PDF generation replacing Documint/Bannerbear, (4) 17 MVP tools in Categories 1-2, and (5) RealizeAI brand design system (#1e3a5f, Garamond/Roboto).

Technology Stack Map



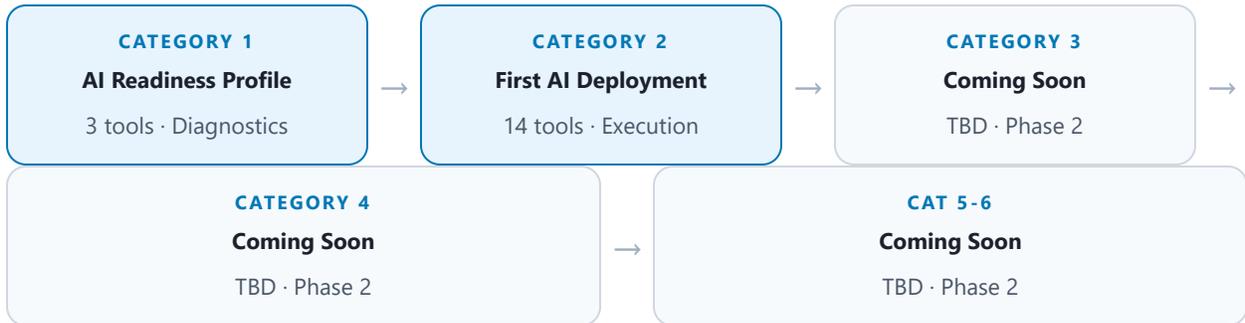
Data Flow Architecture (V2.0 — Form-Before-Download)

User Journey: Persona → Browse → Purchase → Form → PDF → Download

- 1. Persona Mirror** → User self-selects persona (“I Need AI” / “My Pilot is Stuck” / “I’m Asked to Opine”) → recommended category highlighted
- 2. Category Hub** → Cat 1 & 2 active (17 tools grouped by type); Cat 3-6 visible as “Coming Soon”
- 3. Add to Cart** → Single tool (\$97-197), Bundle A (\$397, 5 tools), Bundle B (\$697, all 17), or Annual (\$997/yr)
- 4. Stripe Checkout** → Payment processed → user account created/linked → confirmation email via SendGrid
- 5. Form Required (Non-Negotiable)** → Mandatory assessment form: Org Name, Title, Revenue, Industry, AI Maturity Level, Use Case, Email
- 6. PDF Generation** → Form data + tool_id → OpenAI API generates personalized content (archetype-specific) → puppeteer renders HTML→PDF → stored in S3
- 7. Delivery** → Immediate download OR email delivery via SendGrid → cached PDF available for re-download without re-form
- 8. Benchmarking** → Form data stored anonymized in `form_responses` table → aggregated for Phase 2 benchmarking dashboard

6-Category Hub Implementation (V2.0)

The core user interface features a **Persona Mirror** landing page with 3 self-select personas, leading to a **6-Category Hub**. Categories 1-2 are fully active in MVP; Categories 3-6 are visible but locked with "Coming Soon" badges. Within each active category, tools are grouped by type (Scoring Tools, Templates, User Guides).



Persona Mirror: Three self-select persona cards ("I Know I Need AI But Don't Know How" / "My Pilot is Stuck" / "I'm Asked to Opine") guide users to a recommended category. This is non-gating — users can browse all active categories regardless of persona selection. Persona data is stored for analytics.

Tool Implementation Matrix

Each of the 17 MVP tools falls into one of three implementation categories. Category determines the technical pattern, integrations, and output format.

CATEGORY	PATTERN	INPUT	OUTPUT	COMPLEXITY
Scoring Tool	React form → Supabase → Score Engine → OpenAI API → Puppeteer PDF	6–15 questions + archetype	Personalized multi-page branded PDF	High (3–5 days/tool)
User Guide	React accordion/tabs → Styled content (read-only)	None (downloadable)	Interactive web reference + downloadable PDF	Low (1–2 days/tool)
Template	React form → Template engine → Puppeteer PDF/Word	Structured form fields	Formatted PDF/Word/Excel	Medium (2–4 days/tool)

Category 1: Know Your AI Readiness Profile (3 Tools)

Scoring Tool

Tandem Realization Matrix (TRM)

The moment of truth. 15-min two-axis matrix diagnostic. Plot your organization into one of 4 zones: Tourist, Bureaucrat, Hazard, or Accelerator. 6 dimensions scored on 1–5 scale. Personalized PDF with zone-specific recommendations.

Scoring Tool

3P Alignment Audit (Two-Lock System)

Exposes C-Suite perception gaps. Reinterprets diagnostic scores through Purposeful leadership, Profitable operations, and Pervasive execution lenses. Radar chart + gap analysis PDF.

User Guide

Swamp Draining Reality Check (Data Debt Calculator)

Brutal honesty on legacy data cleanup: how much time and money before you can START a 90-day AI sprint. Cost estimation guide with downloadable framework.

Category 2: Launch Your First AI Deployment (14 Tools)

TOOL	TYPE	EST. DAYS	KEY FEATURES
Master AI Initiative Charter (The Gateway)	Template	3	30+ structured fields → boardroom-ready charter PDF
The 4 Laws Template	Template	2	Governance framework template
Scoping — The 3-Key Vault	Guide	1.5	Interactive scoping methodology guide
The 30 Pearls for a Pilot Project	Guide	1.5	Searchable best-practices catalog
Requirement Gathering Template	Template	2	Structured requirements capture + export
Master AI Risk Register	Template	3	Risk identification framework + Excel/PDF output
CFO Drift Mitigation Planner	Template	3	OpEx drift tracking + mitigation PDF
Macro YoY P&L Model	Template	3.5	Multi-tab financial model Excel/PDF
90-Day Sprint RACI Matrix	Template	2.5	Role-based RACI matrix + exportable PDF
90-Day Implementation Roadmap	Guide	1.5	Interactive sprint-by-sprint roadmap guide
Build vs. Buy Checklist	Guide	1.5	Decision framework with scoring
Procurement Red-Flag Checklist	Guide	1.5	Vendor evaluation red-flag interactive grid
Pilot Purgatory Extraction Audit	Scoring	4	"Kill or Scale" diagnostic with blocker analysis PDF
The Antibody Audit	Scoring	4	12 questions → organizational resistance heatmap PDF

Total MVP Implementation Effort: 20 AI-powered tools — 5 scoring tools + 15 templates & guides. With the agent-augmented model, this is delivered in **5 working days (40 hours)** by one architect + 13 AI agents — see Volume III for the detailed Basis of Estimate.

Core Schema (Supabase / PostgreSQL) — V2.0

TABLE	PURPOSE	RLS POLICY
users	Profile: email, password_hash, name, created_at	<code>auth.uid() = id</code>
categories	6 content categories (name, slug, description, order)	Public read
tools	17 tools mapped to categories (name, type, description, price_min, price_max)	Public read
bundles	Bundle definitions (name, description, price, discount_percent, tools[])	Public read
orders	Purchase records (user_id, bundle_id, total_price, status)	<code>auth.uid() = user_id</code>
order_items	Line items per order (order_id, tool_id, price)	<code>auth.uid() = orders.user_id</code>
form_responses	Mandatory assessment data (org, title, revenue, industry, maturity, use_case, email)	<code>auth.uid() = user_id</code>
pdfs	Generated PDFs (user_id, tool_id, form_response_id, s3_url, regenerated_count)	<code>auth.uid() = user_id</code>
purchases	Stripe payment records, tool access grants	<code>auth.uid() = user_id</code>
user_responses	Assessment inputs per tool per user	<code>auth.uid() = user_id</code>
generated_reports	PDF metadata, S3 URLs, expiration dates	<code>auth.uid() = user_id</code>
tool_feedback	Thumbs up/down + comments per tool	<code>auth.uid() = user_id</code>
benchmarking_data	Anonymized aggregated responses (no PII)	Admin-only read
organization_settings	Enterprise tier: team management, admin controls	<code>auth.uid() = admin_id</code>
organization_members	Multi-seat user mapping for Enterprise tier	<code>auth.uid() IN (admin, member)</code>

Form-Before-Download Access System

Gating is form-based, not credit-based. After Stripe checkout, a Supabase Edge Function handles the access workflow:

1. Verify `orders.status = 'paid'` for user →
2. Present mandatory 7-field assessment form →
3. Validate + insert row in `form_responses` →
4. Trigger OpenAI PDF generation →
5. Store PDF in S3 →
6. Return signed download URL

Form responses power the benchmarking data pipeline (anonymized, aggregated). Users who purchased bundles can access all tools in the bundle without re-purchasing.

T-05

Security Architecture

Security is embedded at every layer — not bolted on post-development. This aligns with BFI's production security track record (2 OWASP sprints, 23 findings remediated) and the BRD's stringent data privacy requirements.



Row-Level Security (RLS)

Aggressive Supabase RLS policies enforce `auth.uid() = user_id` on every data table. User A cannot, under any circumstances, query User B's data. Pen-tested pre-launch.



Authentication

Supabase Auth with magic-link email (passwordless). Session tokens expire after 14 days of inactivity. Future: SSO/SAML for Enterprise Plus tier.



Payment Security

Stripe handles all PCI DSS compliance. BFI never touches raw card data. Stripe.js renders payment fields in secure iframes.



Watermarking

"© RealizeAI Value | Confidential" semi-transparent watermark on every PDF page. Copyright header on every report. 30-day auto-delete from S3.



OWASP Top-10 Audit

Full OWASP security audit at each phase gate. XSS prevention via React's built-in escaping. CSRF protection via Supabase JWT. SQL injection impossible via parameterized queries.



Compliance Stack

GDPR (right to deletion, data export), CCPA compliance, Terms of Service, Privacy Policy, DPA readiness for EU customers, disclaimers on all tools and PDFs.

Third-Party Integrations

SERVICE	PURPOSE	INTEGRATION METHOD	PHASE
Stripe	Payments, billing, subscriptions	Stripe.js + Webhooks to Supabase Edge Functions	1
Supabase	Database, auth, storage, edge functions	Supabase JS client (@supabase/supabase-js)	1
OpenAI API	AI-personalized PDF generation	REST API from Supabase Edge Function	1A
AWS S3	PDF storage with 30-day TTL	S3 SDK via Edge Function; signed URLs for download	1A
CloudFlare	CDN, DDoS protection, static assets	DNS + proxy configuration	1A
SendGrid	Transactional + nurture emails	REST API from Supabase Edge Functions	1B
PostHog	Product analytics, funnels, feature flags	PostHog JS SDK	2
Google Analytics 4	Campaign attribution, UTM tracking	GA4 tag + GTM	1

Performance & Scalability

METRIC	TARGET	APPROACH
Page Load Time	< 2 seconds	CDN-served static assets; code splitting; lazy loading
PDF Generation	< 5 seconds	Pre-compiled templates; optimized chart rendering; async generation with loading state
Lighthouse Score	≥ 90	Image optimization; tree-shaking; minimal JS bundle; efficient CSS
Concurrent Users	1,000+ simultaneous	Supabase connection pooling; CloudFlare caching; stateless frontend
Database Queries	< 100ms p95	Indexed columns; optimized RLS policies; connection pooling
Uptime SLA	99.9%	Supabase managed uptime + CloudFlare failover + Sentry alerting

Scalability Path: The Supabase + CDN architecture scales horizontally without code changes. Moving from 100 to 10,000 users requires only Supabase plan upgrades and CloudFlare cache tuning — zero architectural refactoring.

Sprint Delivery Map

SPRINT	WEEKS	DELIVERABLES	MILESTONE
Sprint 1	1–2	Architecture + Dev Brief V2.0 alignment; Supabase schema (8 tables) + RLS; Stripe integration (singles, bundles, annual); RealizeAI design system; auth flow; CI/CD pipeline	Foundation ready
Sprint 2	3–4	Persona Mirror landing page; Category Hub (6 categories, 2 active); Category detail pages; Form-Before-Download flow; cart/checkout (singles, bundles, annual)	User journey live
Sprint 3	5–6	Cat 1 Tools: TRM (scoring + zone classification + OpenAI PDF pipeline); 3P Alignment Audit (15Q + radar PDF); Swamp Draining Reality Check (guide + framework); legal docs; security audit; CAT 1 LAUNCH	 Category 1 Go-Live
Sprint 4	7–8	Cat 2 Scoring: Pilot Purgatory Extraction Audit; Antibody Audit; Cat 2 Templates: Master AI Initiative Charter; 4 Laws Template; Requirement Gathering Template	5 Cat 2 tools live
Sprint 5	9–10	Cat 2 Templates (cont): Master Risk Register; CFO Drift Planner; Macro YoY P&L Model; 90-Day Sprint RACI; benchmarking data pipeline	9 Cat 2 tools + benchmarking
Sprint 6	11–12	Cat 2 Guides: 3-Key Vault; 30 Pearls; 90-Day Roadmap; Build vs Buy; Procurement Red-Flag; SendGrid email integration; book integration; promo codes	All 14 Cat 2 tools + email
Sprint 7	13–14	Admin analytics dashboard; feedback system; full MVP security audit (17 tools); performance tuning; mobile QA; full production deployment; FULL MVP LAUNCH	 Full MVP Live

BlueFalconInk

Service-Disabled Veteran-Owned Small Business

Volume II: Technical · Build + Operate Proposal for AI Value Realization Accelerator (V4 Hosting-First)

Updated 24 March 2026 · bluefalconink.com

CONFIDENTIAL — For authorized recipients only. © 2026 BlueFalconInk LLC. All rights reserved.