

AI Value Realization Accelerator Platform

Management Volume — Understanding, Approach, Key Personnel, Risk Management, and Project Governance for the RealizeAI Value Micro-SaaS Platform (Developer Brief V2.0). Delivered under a Build + Operate engagement between BlueFalconInk LLC and RealizeAI Value.

Prepared for: Prashant & RealizeAI Value Team

Prepared by: BlueFalconInk LLC

Model: 1 Architect + 13 AI Agents

Pricing: \$5,000 Build + \$249/mo Managed Hosting

Date: 24 March 2026

SDVOSB Certified

14

WEEK TIMELINE

17

MVP TOOLS (CAT 1-2)

5

STAGE PIPELINE

7

SPRINTS



INTERACTIVE COMPANION

Volume I: Management Dashboard

Executive summary, management approach, org chart, risk register, sprint governance, and past performance.

<https://proposal-viewer-42380604425.us-central1.run.app/p/a89dc9963e81194a/vol1>

Executive Summary

BlueFalconInk LLC ("BFI") is pleased to present this Management Volume as part of our **Build + Operate engagement** with RealizeAI Value, in response to the AI Value Realization Accelerator Developer Brief V2.0 (March 2026). We will design, build, and deliver a premium micro-SaaS platform at **realizeaivalue.com** — a form-before-download, product-led growth platform offering mid-market executives AI diagnostic scoring tools, strategic user guides, and execution templates organized into six content categories (MVP: Categories 1-2, 17 tools).

This is a **Build + Operate** engagement: BFI delivers the platform for a **\$5,000 fixed build fee**, then manages ongoing hosting, security, and operations for **\$249/month**. You keep **100% of platform revenue** — no revenue share, no Stripe audits, no monthly exports. BFI's delivery model — one senior architect (**John Eric Lewis**) commanding a **13-agent AI operations fleet** — enables us to deliver the full RealizeAI Value MVP platform within **5 working days (1 week)** at pricing that is 93% below industry standard.

Key Differentiator — Human Led, Agent Driven: BlueFalconInk doesn't just build AI platforms — we *operate* one. Our production fleet of 13 AI agents on 26 Cloud Run services, our security-hardened infrastructure (OWASP-audited, ID-token authenticated), and our veteran-led operational discipline translate directly to faster delivery, fewer surprises, and enterprise-grade quality. The same agent fleet that runs our platform will build and operate yours — with perpetual managed hosting at **\$249/mo** and agents executing at zero marginal cost.

Understanding of Requirements

Client Objective

Build a premium, form-before-download B2B micro-SaaS platform that operationalizes the "AI Value Realization Accelerator" methodology — shifting from traditional consulting lead-magnets to a product-led growth model serving mid-market C-suite executives (CFO, CIO, CEO; \$100M–\$1B revenue companies) who need structured frameworks for AI adoption, governance, and pilot execution.

Alignment Complete (V2.0 → V4 Hosting-First): This proposal is aligned to the updated **Developer Brief V2.0** (March 2026) with **V4 hosting-first pricing**. BFI and RealizeAI Value have completed alignment work during our discussions. The 10 hours typically required for initial requirements analysis have been removed from our estimate.

V2.0 Core Requirements Matrix

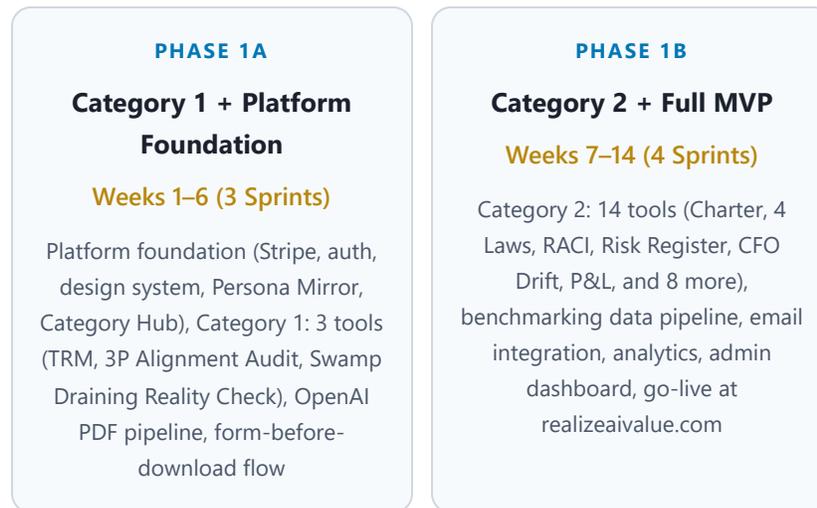
REQUIREMENT AREA	V2.0 SPECIFICATION	BFI UNDERSTANDING
Platform Model	Form-Before-Download, PLG, direct purchase	Stripe checkout → mandatory assessment form → personalized PDF generation and download
Navigation	Persona Mirror → 6-Category Hub (MVP: Cat 1-2)	3 self-select personas → Category Hub with Cat 1-2 active, Cat 3-6 "Coming Soon"
Tool Taxonomy	17 tools (Cat 1: 3, Cat 2: 14) — Scoring Tools, Templates, User Guides	All 17 tools built in MVP; grouped by type within each category
Pricing Architecture	Singles (\$97-197), Bundle A (\$397), Bundle B (\$697), Annual (\$997/yr)	Stripe-integrated cart with single/bundle/subscription options; early-adopter pricing for first 100 users
Frontend Stack	React / TypeScript / Tailwind / shadcn/ui	React SPA with RealizeAI brand design system (Garamond/Roboto, #1e3a5f/#2563eb/#0d9488)
Backend / DB	users, categories, tools, bundles, orders, order_items, form_responses, pdfs	Supabase PostgreSQL with RLS; e-commerce order pattern (no credit engine)
PDF Generation	OpenAI API + puppeteer/wkhtmltopdf; personalized by maturity archetype	4 archetype variants (Tourist/Bureaucrat/Hazard/Accelerator); McKinsey-quality output; cached + retakeable
Benchmarking	Phase 1: collect via mandatory form; Phase 2: premium dashboard upsell (\$297-497/yr)	Anonymized aggregation with min-5-response thresholds; GDPR data deletion support
Book Integration	"The Machine Speaks Human" — chapter references on every tool card	Book branding throughout platform; footer cover image; chapter-to-tool deeplinks
Email / Analytics	SendGrid transactional; PostHog analytics; feedback at multiple touchpoints	Purchase confirmation + PDF delivery emails; NPS, star ratings, abandonment capture
Security / Legal	GDPR/CCPA, RLS, disclaimers	Full compliance stack including TOS, Privacy Policy, DPA readiness, benchmarking opt-in

Scope Confirmation: BFI confirms understanding of the full Developer Brief V2.0 scope across all 13 sections. Our proposal addresses every functional requirement, UX flow, purchasing model, and technical specification documented therein. The V2.0 brief supersedes the original BRD v2.0 — all V2 changes are reflected in this updated proposal.

Delivery Methodology

BFI will execute this engagement under a **Build + Operate model** using an **Agile-Sprint methodology**, with 2-week sprint cycles, continuous stakeholder visibility, and production-first engineering. One senior architect (John Lewis) leads every sprint personally, augmented by a 13-agent AI fleet that executes scaffolding, QA, security, CI/CD, and documentation in parallel. Each sprint concludes with a deployable increment and a live demo. Post-build, BFI transitions to managed hosting operations (\$249/mo).

Two-Phase Delivery Plan



Phase 0 Option Available: A \$6,000 validation sprint is available before committing to the full engagement. Build an AI front-end agent, tool recommender, and payment gate to validate demand. Hours credited toward Phase 1 if you proceed. See Volume III, Section P-13.

V2 Sub-Phase Strategy: The Developer Brief V2.0 expanded MVP scope from 5 to 17 tools. BFI recommends a Category 1 → Category 2 sub-phased approach: launch with 3 Category 1 tools (the diagnostics that validate readiness) then build out all 14 Category 2 tools (the execution frameworks). This preserves the V2 vision while giving Prashant an earlier launchable product and a natural pivot point.

Sprint Cadence & Governance

GOVERNANCE ELEMENT	FREQUENCY	DETAILS
Sprint Planning	Bi-weekly (Mon)	Backlog grooming, sprint goal definition, capacity planning
Daily Standups	Daily (async Slack)	Progress, blockers, and next-actions posted by 9AM CT
Sprint Demo	Bi-weekly (Fri)	Live demo of deployable increment to stakeholders
Sprint Retrospective	Bi-weekly (Fri)	Process improvement; velocity tracking; risk log update
Stakeholder Review	Monthly	Phase gate review, budget tracking, risk assessment, scope alignment
Deliverable Acceptance	Per sprint	Client sign-off on sprint deliverables before proceeding

Communication Plan

Primary Channel	Dedicated Slack workspace with #project, #engineering, #design, #stakeholder channels
Video Meetings	Google Meet / Zoom — Sprint demos, planning, and escalations
Project Tracking	GitHub Projects board with sprint-mapped milestones
Document Repository	Shared Google Drive with versioned deliverables
Escalation Path	John Lewis (BFI) ↔ Prashant (RealizeAI Value) — direct partner channel, no intermediaries

Key Personnel

Organizational Structure — Human Led, Agent Driven



8 of 11 fleet agents are directly assigned to this engagement. All agents report to John — every decision, every code merge, every deployment is human-approved.

Team Profile

John Eric Lewis

Solutions Architect · Lead Engineer · Project Manager — Your Single Point of Contact

Veteran, SDVOSB founder, and full-stack engineer who built and operates a 26-service production AI platform on Google Cloud — the same class of infrastructure your platform requires. John personally leads every engagement: architecture decisions, complex business logic, Stripe/Supabase integrations, code review, stakeholder communications, and quality sign-off. 100% of human hours on this project are John's hours — no handoffs, no junior resources, no coordination overhead.

Web Design Agent

UI/UX Scaffolding & Design Systems

Generates React component structure, Tailwind CSS design systems, responsive layouts, and premium dark-mode aesthetics. Agent-produced scaffolding is reviewed and refined by John for pixel-perfect quality.

SOC Agent (ADKSOC)

Security & Compliance

Executes OWASP Top-10 audits, RLS penetration testing, security scanning, and vulnerability triage. The same agent that secures BFI's production fleet will secure your platform.

NOC Agent

Infrastructure & DevOps

Provisions CI/CD pipelines, deployment automation, monitoring configuration, and infrastructure health checks. Handles CloudFlare CDN, hosting setup, and production environment management.

Legal & Contracting Agents

Compliance Documentation

Draft Terms of Service, Privacy Policy, disclaimers, SOW structures, milestone definitions, and delivery gate verification. Reviewed and finalized by John before client delivery.

Strategic Ops & QA Agents

Planning, Testing & Quality

Assist with sprint planning, progress tracking, test generation, Lighthouse audits, mobile QA, and cross-browser validation. Reports and documentation generated in parallel at zero billable cost.

Why This Model Works: Agents handle ~83% of the work that would traditionally require 4–5 separate specialists — code scaffolding (40%), QA automation (15%), CI/CD + deployment (10%), design system generation (10%), documentation (8%). John focuses exclusively on the 17% requiring human judgment: architecture decisions, complex business logic, integrations, and quality review. Result: 202 human hours delivers what traditionally takes 656 hours and a 5-person team.

Quality Assurance Plan

QUALITY GATE	CRITERIA	FREQUENCY
Code Review	All PRs require 1+ reviewer approval; no direct-to-main pushes	Every PR
Automated Testing	Unit tests (≥80% coverage), integration tests, E2E Playwright tests	Every PR / CI
RLS Pen-Testing	Verify User A cannot access User B data across all tables	Pre-launch + quarterly
OWASP Top-10 Scan	Zero critical/high findings before each phase go-live	Per phase gate
PDF Quality Check	Pixel-perfect rendering, watermarks present, branding correct	Every tool release
Performance Benchmark	PDF generation <5s; page load <2s; Lighthouse score ≥90	Sprint demo
Accessibility Audit	WCAG 2.1 AA compliance on all public-facing pages	Per phase gate
UAT	Stakeholder acceptance testing with sign-off checklist	Per sprint

Risk Management

Risks are actively managed in a living risk register updated each sprint retrospective. The following table maps to the BRD Section 15 risk register with BFI-specific mitigations.

RISK	SEVERITY	PROBABILITY	BFI MITIGATION
Low tool adoption / incomplete assessments	MED	Medium	A/B test tool UX; progress indicators; mobile-first optimization; simplify question flow to reduce dropoff
Low conversion rate (browsers → purchase)	HIGH	Medium	Beta pricing validation with 10–20 users pre-launch; blurred report previews as social proof; progressive disclosure CTAs
Pricing backlash / perceived overpricing	MED	Low	A/B test \$297/\$497/\$697 anchors; "equivalent consulting: \$35,000+" framing; Tier 1 impulse pricing (\$49–\$97)
Data breach or security incident	HIGH	Low	Supabase managed security + aggressive RLS; pre-launch pen-test; OWASP audit; E&O / cyber insurance; zero PII in benchmarking data
OpenAI API cost/latency at scale	MED	Medium	Cache PDFs aggressively in S3; rate-limit retakes (1/month free); monitor API cost per PDF (\$0.50-2.00); fallback to template-only PDF if API down
Lovable.dev dependency / availability	LOW	Low	Standard React/TypeScript underpins Lovable; fallback to custom React if needed; zero vendor lock-in on component layer
17-tool MVP scope overrun	MED	Medium	Sub-phase into Cat 1 (3 tools, Week 6 gate) then Cat 2 (14 tools, Week 14); strict sprint-level scope locks; change request process
Book launch delays impacting integration	LOW	Medium	Platform functions independently per BRD; book hub is MVP placeholder; QR integration is Phase 2 and can be deferred

Relevant Past Performance

BLUEFALCONINK PRODUCTION AI PLATFORM

18-Service Cloud Run Fleet — Live Production

Designed, built, secured, and currently operates a production AI platform on Google Cloud with **26 Cloud Run services** (13 ADK-powered agents, 13 supporting services). Completed two OWASP security sprints, implemented ID-token authenticated cross-agent communication, and maintains <\$58/month infrastructure costs serving real workloads.

26

CLOUD RUN SERVICES

13

AI AGENTS

<\$58

MONTHLY INFRA COST

100%

UPTIME (Q1 2026)

RELEVANT TECHNOLOGY ALIGNMENT

Direct Stack Overlap with BRD Requirements

BRD REQUIREMENT	BFI PRODUCTION EXPERIENCE
React / TypeScript frontend	Next.js + React production apps (bluefalconink.com)
PostgreSQL / Supabase backend	Cloud SQL + Firestore production databases
Stripe payments integration	Payment processing integration experience
Premium PDF generation	Dynamic report generation pipeline (investor prospectus, capabilities statement)
Cloud hosting with auto-scaling	Google Cloud Run with auto-scaling, load balancing, CDN
OWASP security hardening	2 completed security sprints; 23 findings remediated; zero critical open
CI/CD deployment pipelines	Cloud Build + Artifact Registry automated pipeline
Analytics / monitoring	Structured logging, Cloud Monitoring, custom dashboards

SDVOSB + Build & Operate Advantage: BlueFalconInk LLC is a certified Service-Disabled Veteran-Owned Small Business. This engagement operates as a Build + Operate agreement with perpetual managed hosting — meaning BFI is invested in the platform's long-term success through ongoing operations, not just delivering a one-time project. You keep 100% of platform revenue. Engagements with BFI may also qualify for set-aside procurement preferences, SBA mentor-protégé programs, and government subcontracting credits where applicable.

M-08

Assumptions & Dependencies

ASSUMPTIONS

1. This engagement operates under a Build + Operate agreement: \$5,000 fixed build + \$249/mo managed hosting. Client retains 100% of platform revenue (see V4 BOE for details).
2. Client will provide timely feedback within 3 business days on sprint deliverables.
3. Developer Brief V2.0 (March 2026) represents the approved scope baseline; changes follow formal CR process.
4. Client owns the realizeaivalue.com domain and Stripe account (or BFI sets up under client ownership).
5. Tool question sets, scoring algorithms, and framework methodologies will be provided by client or developed collaboratively.
6. "The Machine Speaks Human" book content/QR code mappings provided by client when ready.
7. Third-party services (Supabase, Stripe, PDF API) maintain published SLAs during engagement.
8. BRD alignment work is already complete — those hours are excluded from the estimate.

DEPENDENCIES

1. **Lovable.dev platform stability** — mitigated by fallback to standard React.
2. **Supabase availability** — 99.9% SLA; mitigated by managed service guarantees.
3. **Stripe PCI compliance** — maintained by Stripe; BFI does not handle raw card data.
4. **OpenAI API account** — client-owned API key for PDF personalization; BFI configures prompt engineering.
5. **Legal review** — TOS/Privacy Policy generated via termly.io; optional attorney review at client's discretion.

BlueFalconInk

Service-Disabled Veteran-Owned Small Business

Volume I: Management · 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.