

Product Roadmap 2026 — 2027

AI Portfolio Rationalization Advisor · 24-Month Plan

Prepared by K. Turner, Senior AI Product Manager

Version 1.0 · May 2026 · Confidential

Phase	Name	Period	Target	Focus
1	Validate	Months 1–4	3 design partners	Problem-solution fit
2	Launch	Months 5–10	10 paying customers	Product-market fit
3	Scale	Months 11–18	50 customers	Repeatable growth
4	Expand	Months 19–24	200–500 customers	Platform & exit

24-Month Product Roadmap

Phase 1	Validate	Months 1–4
<i>Find 3 design partners. Prove the problem is real. Build only what's needed to learn.</i>		
Product	Add user auth (Supabase). Multi-user sessions. Decision persistence across sessions.	AI / ML
Go-to-market	Recruit 3 design partners (free access). Weekly feedback sessions. Define ICP — who feels this pain most acutely?	Platform
Target: 3 design partners	Revenue: \$0 (free beta)	Key question: Will they pay?

Phase 2	Launch	Months 5–10
<i>Convert design partners to paying customers. Ship the AI core. Find the repeatable motion.</i>		
Product	Self-serve onboarding flow. Tiered pricing live (Starter / Growth / Enterprise). PDF cycle reports. Email notifications for pending decisions.	AI / ML
Go-to-market	ServiceNow partner ecosystem outreach. Content marketing (LinkedIn, ITAM publications). First 10 paying customers. Case study from design partner.	Platform
Target: 10 customers	Revenue: \$5K–15K MRR	Key question: Can we retain them?

Phase 3	Scale	Months 11–18
<i>Grow to 50 customers without being present in every deal. Build the agentic AI layer.</i>		
Product	Mobile-responsive UI. Advanced governance workflows. Custom signal weighting per customer. Full API access for enterprise tier.	AI / ML
Go-to-market		Platform

Go-to-market

Channel partnerships with ServiceNow SIs. Marketplace listing. Published case studies. Seed raise or strategic acquirer conversation initiated.

Platform

SOC 2 Type II certified. SSO / SAML for enterprise. Azure DevOps integration. LangSmith for AI observability and prompt management.

Target: 50 customers	Revenue: \$50K–100K MRR	Key question: Can we grow without me?
----------------------	-------------------------	---------------------------------------

Phase 4 Expand Months 19–24

500-customer architecture. Adjacent markets. Platform play or strategic exit.

Product

Full agentic cycle management. AI-generated board-level reports. Scenario modeling (what if we retired these 10 apps?). Signal marketplace for custom integrations.

AI / ML

Proprietary model trained on anonymized customer portfolio data. Industry benchmarking. Predictive cost and risk modeling. Multi-model ensemble for recommendations.

Go-to-market

Enterprise direct sales team. Expand to adjacent markets: cloud cost governance, vendor rationalization, tech debt management. Series A or acquisition.

Platform

Multi-region deployment (US, EU). Enterprise SLAs. White-label option for SI partners. Full compliance audit trail for regulated industries.

Target: 200–500 customers	Revenue: \$500K+ MRR	Key question: Platform or exit?
---------------------------	----------------------	---------------------------------

Strategic Notes

Why phased scaling?

Building for 500 customers before finding product-market fit is the most common way technically capable founders kill products. Each phase gate is a deliberate checkpoint — only invest in the next phase when the current phase's key question is answered with paying customers, not assumptions.

The Centaur model as moat

The human-in-the-loop governance framework is not just a feature — it's a defensible architectural position. Enterprise AI buyers are under pressure to demonstrate AI accountability and auditability. The Centaur model delivers this out of the box. Competitors adding AI to existing portfolio tools will struggle to retrofit this.

The ServiceNow bet

ServiceNow's APM module is expensive and complex. Their mid-market customers — companies with 50–500 applications — are underserved. A lightweight, AI-native alternative that integrates with ServiceNow rather than replacing it is a wedge, not a threat. This positions the product for acquisition by ServiceNow or a major SI partner.

AI cost management at scale

At 500 customers making 500 queries/month, Anthropic API costs become material. Phase 2 introduces model routing to minimize cost. Phase 3 introduces a fine-tuned model on customer data that reduces dependence on frontier models for routine queries. Phase 4 explores an ensemble approach with a proprietary base model.

Build vs. buy decisions

Auth (Supabase), vector store (pgvector on existing Postgres), observability (LangSmith), and compliance (Vanta for SOC 2) are all buy decisions. The product team should only build what is proprietary and differentiating — the governance model, the signal engine, and the customer-facing product experience.

This roadmap is confidential and intended for interview, investor, and strategic partner use only. Prepared by K. Turner, Senior AI Product Manager, May 2026.