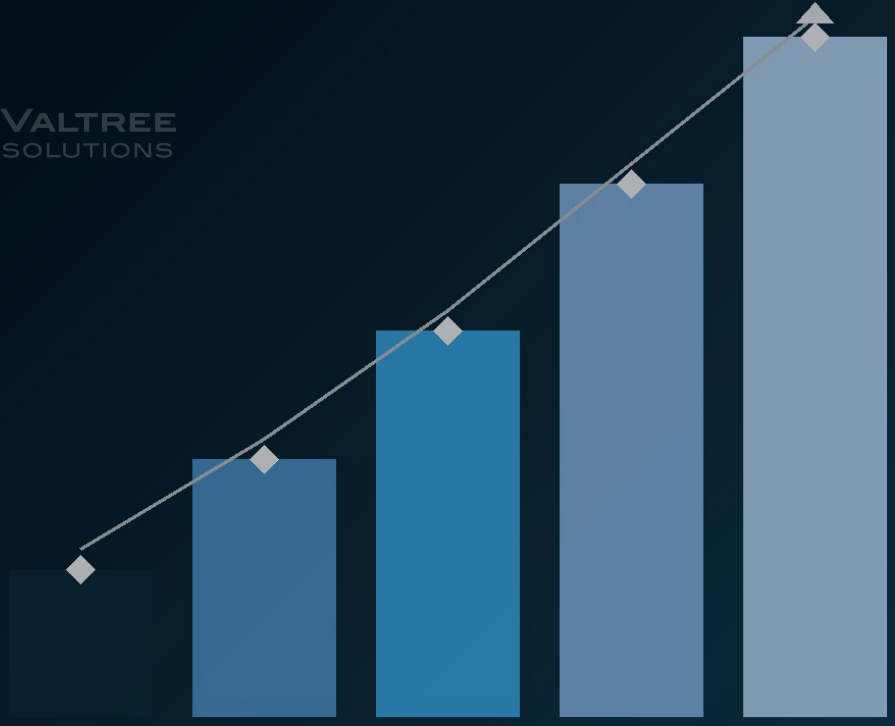
VALTREE SOLUTIONS



VALTREE PRODUCT MANAGEMENT WHITE PAPER SERIES

Applied AI Transformation & Enterprise PMO

In-House AI Integration, Buy-vs-Build Governance & Value Capture

$18M	45,000+	100%
ANNUAL COST SAVINGS	ACTIVE INTERNAL USERS	DATA LEAKAGE PREVENTION

Executive Summary

This paper describes how Valtree's embedded product leaders led an enterprise-wide AI transformation for a global financial institution, establishing AI governance, a buy-vs-build decision framework, and a bespoke internal productivity assistant used across the organization. The client entered the engagement with fragmented, uncoordinated AI experiments spread across dozens of business units — some sanctioned, many not — creating both compliance exposure and tens of millions of dollars in redundant vendor spend.

Valtree's team built a risk-based governance framework, a standardized decision matrix for when to buy third-party AI tools versus build proprietary ones, and led product management for a secure, enterprise-wide internal assistant connected to the firm's proprietary document stores. Within a year, the program delivered eighteen million dollars in annual cost savings, reached more than forty-five thousand active internal users, and recorded zero incidents of data leakage across the assistant's entire usage base — a threshold the firm's risk committee had treated as a precondition for continued investment.

1. Strategic Engagement Context

A Fortune 100 financial services firm had, like many large enterprises in 2024 and 2025, watched AI experimentation spread organically across its business units faster than governance could keep pace. Individual teams stood up their own pilots with whichever vendor a team lead had encountered at a conference, with no central visibility into what data those tools touched or what regulatory exposure they created.

Two consequences had become impossible to ignore by the time Valtree was engaged. Unregulated shadow AI usage — employees routing sensitive client and firm data through consumer-grade AI tools without sanction — posed a genuine compliance and data-privacy risk that internal audit had begun flagging formally. And overlapping, uncoordinated vendor purchases across business units were driving tens of millions of dollars in redundant licensing costs for capability the firm frequently already owned elsewhere in the organization.

The firm's leadership wanted a way to say yes to AI adoption faster, not slower — but only if it could be done inside a governance model that satisfied the risk committee, and only if spend could be consolidated enough to fund a genuinely enterprise-grade internal assistant rather than dozens of disconnected point solutions.

2. Core Interventions & Areas of Excellence

Enterprise AI Governance Framework. We formulated a risk-based triage model that evaluated every AI initiative in the portfolio against privacy, regulatory compliance, and security vectors, sorting initiatives into fast-track, standard-review, and restricted tiers. This let genuinely low-risk experimentation move quickly while concentrating compliance review effort on the initiatives that actually warranted it.

Buy vs. Build Decision Matrix. We established standardized architectural criteria — covering differentiation, data sensitivity, total cost of ownership, and time-to-value — to determine when the firm should procure a third-party AI solution versus invest in a proprietary internal capability. This matrix became the mandatory first step in the firm's AI initiative intake process, immediately surfacing overlapping purchases that could be consolidated.

Internal AI Assistant Deployment. Valtree led product management end-to-end for an enterprise-wide internal conversational assistant, securely connected to the firm's proprietary document stores through a private-tenant architecture with PII redaction built into the retrieval layer. We managed the assistant's roadmap, rollout sequencing across business units, and the internal change-management program that drove adoption.

Centralized AI Steering Committee. To make the governance framework and decision matrix stick organizationally, we helped stand up a Centralized AI Steering Committee with representation from risk, technology, and the business units, replacing the prior pattern of decentralized, team-level purchasing decisions.

3. Organizational Maturity Matrix

Pillar	Ad-Hoc State (Pre-Intervention)	Managed Enterprise State
Vendor Procurement	Decentralized team purchases	Centralized AI Steering Committee approval
Data Privacy	Unfiltered public API calls	Private tenant routing with PII redaction
ROI Attribution	Anecdotal / unmeasured	Automated productivity & unit cost tracking
Risk Classification	Inconsistent, ad hoc review	Standardized fast-track / standard / restricted tiers
Internal Tooling	Fragmented point solutions	Single enterprise assistant, firm-wide

4. Implementation Roadmap

The transformation program ran across four phases over twelve months. Phase one established the governance framework and steering committee, giving the firm a mechanism to evaluate the dozens of existing shadow-AI initiatives already in flight. Phase two applied the buy-vs-build matrix retroactively across the existing vendor portfolio, identifying consolidation opportunities that funded much of the subsequent build work. Phase three designed and built the internal assistant's core architecture, prioritizing the private-tenant routing and PII redaction layer before any business unit was onboarded. Phase four rolled the assistant out business unit by business unit, pairing each rollout with a change-management push led jointly by Valtree and the client's internal communications team.

5. Results & Impact

Twelve months after program kickoff, the consolidated governance and procurement model was delivering eighteen million dollars in annual cost savings, primarily through eliminated redundant vendor contracts and improved negotiating leverage from centralized purchasing. The internal assistant reached more than forty-five thousand active users across every major business unit, and — critically for a regulated financial institution — recorded zero data-leakage incidents across its entire usage base, satisfying the precondition the risk committee had set before approving firm-wide rollout.

6. Key Learnings

The clearest lesson from this engagement is that governance and velocity are not in tension when the framework is designed around risk-based triage rather than blanket review — most AI initiatives in a large enterprise are genuinely low-risk and should move quickly, freeing compliance capacity for the small number that need real scrutiny. Second, a buy-vs-build matrix pays for itself before a single new feature ships, simply by surfacing redundant spend that already exists across business units. Third, technical safeguards like PII redaction and private-tenant routing are necessary but not sufficient — durable trust with a risk committee also requires visible, ongoing measurement, not a one-time architecture review.

7. Recommendations for Enterprise & PMO Leaders

Stand up a risk-based governance framework before your organization's shadow AI usage becomes large enough to be a headline risk — retrofitting governance onto years of ungoverned adoption is far more disruptive than building it in early. Require a buy-vs-build assessment as a mandatory gate for any new AI spend, since redundant purchasing is one of the most reliable and fastest sources of funding for a genuine internal capability. And treat your internal assistant's data-privacy architecture as the product's core feature, not an implementation detail — it is what will ultimately determine how far your risk committee lets you scale.

Conclusion

Enterprise AI transformation succeeds or stalls on governance as much as on technology. Valtree's Applied AI Transformation practice pairs product leadership with the governance and change-management discipline that large, regulated organizations need to move from fragmented experimentation to a measurable, trusted, enterprise-wide AI capability.

About Valtree AI

Valtree is a specialized AI product management practice embedding senior product leaders directly inside frontier research labs, developer platforms, vertical SaaS vendors, and global enterprises. We pair deep technical fluency with product and governance discipline to turn AI initiatives into measurable, trusted, production-grade outcomes. Learn more at valtree.com.