DEPLOY, SECURE, AND ORCHESTRATE AGENTIC AI

MILL POND RESEARCH

THE PROBLEM WE SOLVE

Organizations eager to harness agentic AI face a frustrating choice: move fast with fragmented tools that create data silos and vendor lock-in, or move slowly through endless pilots that never reach production. When they switch models, institutional knowledge is lost and teams must start from scratch. Meanwhile, employees across every department use AI tools that operate outside IT's visibility—creating data leakage and compliance risks that traditional cybersecurity frameworks cannot address. CIOs and CISOs find themselves trapped in a zero-sum trade-off between enabling innovation with unmanaged AI that creates risks they cannot see, or locking everything down with governance that kills productivity.

THE MILL POND RESEARCH SOLUTION

Mill Pond Research provides the industry's only end-to-end platform that enables enterprises to deploy, secure, and orchestrate agentic AI at scale—lowering AI risk while maximizing operational efficiency.

With Mill Pond Research, security and innovation stop being opposing forces. Organizations that build robust security infrastructure into their AI deployment can actually **accelerate** innovation—addressing the legitimate concerns that create organizational hesitation while unlocking agentic AI's full productivity potential.

WorkBench is the premier platform for authoring, testing, and deploying AI agents that become lasting institutional assets. Its model-agnostic architecture preserves organizational context regardless of underlying LLM, eliminating vendor lock-in while professional-grade tools accelerate the path from prototype to production.

Xilos is the governance layer for agentic AI—intelligent middleware that observes, secures, and orchestrates all AI activity across your network. It intercepts and analyzes every prompt, enforces data protection policies in real-time, and maintains comprehensive audit trails for compliance.

Together, they create the secure foundation for confident, enterprise-scale AI implementation.

"While building protocols to audit AI data, we have seen under the hood of many orchestration technologies. Mill Pond Research's insight into the security implications of AI are well beyond anything we have seen so far."

— Jim Fournier, CEO, JLINC

WHY CHOOSE US

- **The Only End-to-End Platform:** While competitors offer point solutions for either AI development or AI security, Mill Pond Research delivers both in a unified platform. This integrated approach eliminates the complexity of stitching together disparate tools and ensures that security is embedded into the AI lifecycle from day one—not bolted on as an afterthought.
- **True Model Independence:** Our model-agnostic architecture means your investment in institutional knowledge, custom agents, and workflows is never held hostage by a single vendor. As the AI landscape evolves—and it will—you maintain the flexibility to adopt new models without sacrificing what you've built.
- **Complete Visibility and Control:** Unlike legacy security tools designed for a pre-AI world, Xilos was purpose-built to address the unique challenges of autonomous systems. Every prompt, every response, every AI interaction across your network becomes visible, governable, and auditable.
- **Deployment on Your Terms:** Whether your security posture demands managed SaaS, private cloud, or fully air-gapped on-premises deployment, our platform adapts to your requirements—not the other way around.

MILL POND RESEARCH

YEAR FOUNDED
2023

FINANCING ROUND
Seed

STRATEGIC PARTNERS
- US AI Safety Institute Consortium
- NIST
- US Department of War
- IBM
- Oracle
- SailPoint

MARKET CATEGORY
- AI Security
- AI Governance

FOUNDERS
Andrew Shimshock
Co-founder & CTO

Pete Shimshock
Co-founder & CAIO

WEBSITE
millpondresearch.com

WorkBench

Author, Test, and Deploy Agentic AI

WorkBench is the premier platform for building and deploying AI agents that become lasting institutional assets:

- **Multi-model access.** Unified interface to state-of-the-art models from OpenAI, Anthropic, Google, and open-source providers
- **Advanced agent creation.** Granular control over model assignment, parameters, plugins, and knowledge base access
- **Knowledge base integration.** Deploy proprietary organizational data for superior, contextually relevant outputs

Xilos

The Governance Layer for Agentic AI

Xilos operates as intelligent middleware that observes, secures, and orchestrates all AI activity across your network:

- **Intelligent intent-based routing.** Dynamic query routing based on semantic analysis, sensitivity, and required capabilities
- **Model-agnostic integration.** Unified API surface supporting public providers and privately hosted models
- **Proactive Data loss prevention.** Automated masking or blocking for real-time detection of PII, credentials, and proprietary data