



Put AI agents to work with institutional control.

AIOS is the control plane for governed agentic work, connecting agent management, workflows, execution environments, governance, and evaluation in one operating system.



The Problem

AI agents are moving into real work faster than the controls around them. As agent work crosses models, tools, systems and teams, institutions lack a consistent way to govern identity, authority, evidence, cost, and human oversight.



The Solution


Zafin AIOS is an operating system for agentic workflows in regulated industries. It gives institutions a governed way to put AI agents to work. The platform combines workflow orchestration, agent management, governance, observability with human authority, evaluations, and cost controls into a model-agnostic layer that sits above AI models and tools.

Proven first inside Zafin across real agentic software delivery and operational workflows.

One Governed Operating Path

Five connected product areas. One governed operating path.


AIOS connects the capabilities institutions need to set agents up, manage their work, run them in approved environments, govern every outcome and improve performance over time.



AGENT MANAGEMENT

Set every agent up with a clear role, access and accountability.

Treat each agent as a governed workforce member, with a defined identity, profile, knowledge, skills and tools.




- **Identity and access:** Give each agent a secure identity and role-based access.
- **Models:** Centrally manage approved providers and route work to the right model for the task based on performance, policy and cost.
- **Skills Registry:** Capture reusable know-how and assign it to the right agent profiles.
- **Knowledge Fabric:** Give agents the approved documents, code, policies and prior case studies they need to begin.
- **Tools Registry:** Connect agents to the right MCP-based tools so they can act in approved systems.

WORK MANAGEMENT

Manage the full lifecycle of agent work.

Turn requests into governed cases, guide them through workflows, and use reviewed results to improve future work.




- **Case types:** Create one-time, scheduled or continuous work.
- **Case setup:** Start each case with the right context, routing, and compute.
- **Workflows:** Guide work through preconfigured or custom steps.
- **Human gates:** Add review, approval, feedback and questions where required.
- **Proof of work:** Retain a complete, auditable record of decisions, actions and results.
- **Knowledge additions:** Turn reviewed work and feedback into reusable case studies.

EXECUTION ENVIRONMENT

Run agent work in an approved environment.

Prepare the physical or cloud environment where agent work runs, with identity, access, and resources ready before execution begins.



- **Physical compute:** Turn approved devices into secure agent environments.
- **Cloud compute:** Configure virtual machines or containers around each workload's access, capacity and runtime needs.

CONTROL TOWER

Monitor, govern and prove every agent outcome.

Bring traces, analytics, cost, access, and proof of work together in the Control Tower.




- **Traces:** See the prompt, output, model, token use and execution time for every turn.
- **Analytics:** Track outcomes, rework, repeated mistakes and human input.
- **Cost:** Attribute model and compute costs to every case.
- **Role-based access:** Control who can view and act on agent work.
- **Proof of work:** Retain the plans, transcripts, approvals, artifacts and test evidence behind each outcome.

EVALUATIONS

Measure how agents perform and where to improve.

Evaluate cases across their lifecycle using rubrics, benchmark sets, and experiments built around your own institution's work.



- **Rubrics:** Define quality with deterministic signals, LLM-as-judge criteria and weighted scores.
- **Benchmark sets:** Test agents against historical cases that reflect real institutional work.
- **Lifecycle evaluations:** Measure performance during a workflow, on demand or on a schedule.
- **Scores and reasoning:** See performance, decision reasoning and where people need to step in.
- **Experiments:** Compare models on your own cases, data and environment.

Ready to take control of agentic work?

Request a briefing with the Zafin AIOS team to discuss your agentic AI, workflow or modernization priorities, define a high-value starting point for your organization, and identify where agentic work can create measurable value with authority, cost controls, and proof of work built-in.

[Book a briefing](#)

