

XTAI Orchestrator

Scale AI with Greater Control and Cost Visibility.

Centralized governance, practical guardrails, and actionable insight into AI consumption.

RECOGNIZED IN THE MARKET



Gartner predicts more than 40% of agentic AI projects will be canceled by the end of 2027, driven by escalating costs, unclear business value, or inadequate risk controls (Gartner, June 2025). The failure point usually isn't the model — it's the gap between a working demo and something a CISO will approve, an auditor will accept, and a CFO can forecast. XTAI Orchestrator, including the agent management capabilities built into it, is designed to close that gap.

What It Is

XTAI Orchestrator is the unified command layer for enterprise AI operations: a single governed view across AI agents, managed infrastructure, cloud environments, and business services.

Built into Orchestrator are its agent management capabilities: the control plane every agent runs through before it ever touches your environment. They decide what an agent may do, verify it hasn't been tampered with, enforce who approved it, and can pull it from production in minutes.

Orchestrator is also where you see the result — what every agent did, what it cost, whether it stayed inside policy, and how it's performing against your SLA — without a ticket or a monthly report. One set of capabilities governs, the other proves, both inside the same platform.

What this looks like: the Orchestrator dashboard — FinOps spend, top cost drivers, open anomalies, agent runs and model mix in one always-current view, the picture a CFO and a platform owner can both read.



Security | Compliance | Guardrails | FinOps



Govern and control AI agents throughout their lifecycle.

- Live observability and complete action history
- Governed knowledge sources and access
- Runtime policies and intervention



Adopt, configure, and deploy enterprise AI agents.

- XTAI Voice, XTAI Headstart, and XTAI Frontline
- No-code configuration aligned to customer needs
- Custom agents for specific workflows



Apply AI-driven intelligence across managed networks.

- Network health and performance visibility
- Event correlation and anomaly detection
- AI-assisted investigation and automation



Govern and operate cloud and desktop workloads.

- Workload and service visibility
- Provisioning, governance, and automation
- Consumption and optimization insights

Model & Platform Agnostic

Support multiple models, providers, third-party AI services, tools, and integrations without rebuilding the governance layer.

Security | Compliance | Guardrails | FinOps

Agent Management Capabilities — Within XTAI Orchestrator

Feature	What it Delivers
Solutions Library <i>agents and workflows that already exist</i>	Prebuilt agents — triage, hydration, voice and knowledge retrieval, ready on day one Fast to stand up — configuration against your systems, not a data-engineering build Proven in production — in use across XTIUM's own operations for nearly two years, automating a large majority of its service desk workflow volume
Governed Agents Lifecycle <i>nothing reaches production unreviewed</i>	Second-person approval — the author of an agent can't be the one who approves it Versioned and reversible — a change that misbehaves is rolled back, not debugged live Emergency withdrawal — a suspect agent is pulled across every environment in minutes
Policy Control <i>rules enforced before work happens</i>	Three levels — platform, customer and agent rules combine into one ruleset — each can only narrow, never widen Spend stopped in path — over-budget work is denied before it's billed PII/PHI protection — sensitive data is redacted automatically, not configured agent by agent
Agent Identity <i>every agent is a named principal, not a shared account</i>	Attributable actions — every action traces to a specific agent and version Instant revocation — access is revoked the moment an agent is removed Role-based access — who can view, approve or run an agent is controlled centrally, not per tool
Model Freedom <i>no single-vendor lock-in</i>	Routed on fit — cost, latency or sensitivity decides which model runs, across more than 40 supported models Automatic fallback — a provider outage degrades gracefully rather than stopping work Deploy where it must run — on-premises and sovereign options for workloads that can't leave your environment
Audit & Traceability <i>the record a security review actually asks for</i>	Tamper-evident — records can't be altered, not even by an admin Decision-level detail — the why behind a decision, not just the what
Open Ecosystem <i>adds to your estate rather than replacing it</i>	Certified partner agents — third-party agents pass the same review and signing bar as XTAI's own Callable by your agents — agents you've already built can invoke XTAI agents as a governed tool


XTAI Orchestrator Capabilities — Proves Every Result

Feature	What it Delivers
Operational Intelligence <i>one view across an estate that spans many tools</i>	Cross-domain in one place — no monthly assembly of separate reports done by hand Always current — reflects the environment now, not at last reporting cycle
Cost Intelligence & FinOps <i>every dollar attributed to the agent that caused it</i>	No surprise invoices — spend visible per agent as it happens, not at month end Cost per outcome — reported against work completed, connected to business results and departmental usage

Feature	What it Delivers
Predictive Analytics <i>see what's coming, not just what happened</i>	Cost trajectory — month-end spend projected against trend, with variance flagged before it lands Capacity and risk signals — constraints and incident-precursor patterns surfaced while there's time to act
Guardrails & Anomaly Detection <i>risk and cost outliers caught and explained</i>	Root-cause attribution — a cost or latency spike is linked to the specific agent and model behind it Third-party guardrails included — native XTAI controls plus Lakera Guard, Azure Content Safety and OpenAI Guardrails in one view Blocks, masks and flags PII and prompt-injection attempts — in real time, not after the fact
Compliance Evidence <i>audit material generated continuously</i>	Mapped to frameworks — aligned to NIST AI RMF and ISO/IEC 42001 Exportable evidence — usable in the customer's own audit, exported as CSV, JSON or PDF
Agent Visibility <i>see what the agents did without asking</i>	Run history — every execution with outcome, timing and cost No ticket required — answers available directly, with a link back to the ServiceNow/JIRA record
Service Assurance <i>proof the service is performing as contracted</i>	SLA against target — attainment reported with the method stated openly Proactive alerting — notified of conditions that matter before they become problems
Natural Language Query <i>ask the platform, don't dig for it</i>	Plain-language questions — “why did spend jump this week” or “show me open high-risk approvals”, answered without a custom report Automated executive briefings — AI-generated summaries of what changed and what needs attention, without manual report creation Action from the answer — stage an approval or an investigation without switching screens

Where This Shows Up

The Guardrails view inside Orchestrator's agent management capabilities, with Natural Language Query alongside it: ask “why did spend jump this week” or “show me open high-risk approvals” in plain language and get an answer or stage an action, without leaving the chat.



About XTIUM

XTIUM is a global AI Native MSP trusted by more than 1,700 organizations worldwide. Through its XTAI platform, XTIUM helps businesses assess, govern, deploy, and operate AI while delivering secure cloud, network, cybersecurity, digital workplace, and communications services under a single accountable partnership. XTIUM is a four-time Gartner® Magic Quadrant™ Leader in Managed Network Services, recognized as the #1 provider for Automation & AI in Gartner's 2026 Critical Capabilities report, and recipient of the 2026 AI Break-through Award for AIOps Solution of the Year.