VENTRALIS

ENTERPRISE PILOT PROGRAM | AUTONOMOUS AI GOVERNANCE

Govern what AI agents can do. Verify what actually happened.

VENTRALIS is an experimental runtime governance platform designed to evaluate, constrain, and audit autonomous AI-agent actions. Our paid pilots help organizations test enforcement boundaries against their own approved workflows in controlled, non-production environments.

THREE WAYS TO ENGAGE

PACKAGE	PROPOSED FEE	DURATION	BEST FOR
Technical Validation	$35,000–$50,000	4–6 weeks	One focused workflow; initial fit and evidence
Enterprise Governance FLAGSHIP	**$125,000**	12 weeks	Up to two workflows; deeper integration and validation
Strategic Co-Development	$200,000–$300,000+	12–16+ weeks	Custom engineering, extended evaluation and roadmap

WHAT WE EVALUATE

Unauthorized execution: Can defined prohibited actions be denied before downstream side effects?

Agent manipulation: How do agents respond to prompt injection, tool-call hijacking, permission misuse, and goal changes?

Operational fit: What are the effects on authorized workflows, false positives, latency, resilience, and evidence completeness?

APPLICABLE ENVIRONMENTS

Financial services and digital assets • Frontier AI labs • Healthcare administration • Pharmaceutical research • Energy and utilities • Enterprise agent platforms

FLAGSHIP PILOT: 12-WEEK DELIVERY PLAN

Weeks 1–2 | Discovery, threat modeling and approved test plan

Weeks 3–4 | Sandbox integration and policy baseline

Weeks 5–7 | Pre-registered benign and adversarial test battery

Weeks 8–9 | Reliability, latency and failure-mode evaluation

Weeks 10–11 | Customer reproduction review and limitations register

Week 12 | Executive findings, technical evidence and next-phase roadmap

DELIVERABLES & ACCEPTANCE

The customer receives an approved test plan, configuration summary, versioned run records, allowed/denied action evidence, ground-truth metrics (TP/FP/TN/FN), latency results, a documented limitations register, and a final executive and technical report.

Acceptance is based on delivery of agreed evidence and reproducible methods—not a guarantee of perfect detection. Critical-action containment and fail-closed behavior are evaluated against mutually approved test cases; unsuccessful tests remain reportable.

COMMERCIAL FRAMEWORK

Flagship payment schedule: 25% at kickoff ($31,250); 25% at integration acceptance ($31,250); 30% on validation and reliability evidence delivery ($37,500); 20% on final report delivery ($25,000). Other tiers use individually negotiated milestones.

Customer responsibilities: Assign engineering and security sponsors; provide approved sandbox access, synthetic test data, relevant APIs and logs, and timely review decisions.

Explicit exclusions: Production access, real-money transactions, live custody or clinical control, unapproved penetration testing, regulatory certification, 24/7 incident response, unlimited integrations, and guarantees against all attacks. Infrastructure and third-party costs are scoped separately.

Intellectual property: VENTRALIS retains its pre-existing technology and reusable methods; customer data and pre-existing IP remain the customer's. Licensing, custom work, confidentiality, liability, and security terms require a separately executed agreement.

NEXT STEP

Schedule a discovery conversation to identify one or two high-risk agent workflows, select a pilot tier, and establish mutually agreed technical and commercial acceptance criteria.

VENTRALIS™ | Powered by MoodJi

Confidential proposal overview • November 2026 • Indicative pricing, subject to scoping and contract • No existing customer relationship implied. Research outcomes are preliminary and require customer-specific validation.