



AI GOVERNANCE & FIDUCIARY SAFETY

Agents, Human Command, Fiduciary Duty Preserved

You are handing AI agents your advisor data, CRM records, and outbound communications. This explains how Alpheous keeps humans in command through structured automation, SEC-aware compliance, and enforceable governance.

alpheous.com

Prepared for evaluation. © 2026 Alpheous, Inc.

Contents

01	What Are Structured Agents	
02	The Approval Framework	
03	Compliance & Governance	
04	Agent Boundaries	
05	Trust Escalation	
06	Structured in Practice	
07	Distribution Governance	

Not Autonomous AI. Structured Agents.

Alpheous agents are Structured Agents: deterministic code producing identical results every time, with AI invoked only where human judgment is required.

Alpheous runs structured automation on dedicated infrastructure: API integrations, rule-based pipelines, scheduled jobs, and compliance checks. AI is the exception, not the rule. This distinction is critical for fiduciary duty; your firm needs predictable, auditable operations.

Deterministic by Default

Advisor data processing uses structured rules: territory matching, advisor classification, document routing, and compliance flagging. AI weighs in only on ambiguous cases requiring judgment.

AI for Judgment Only

AI drafts market commentary in your firm's voice, synthesizes research data, and generates quarterly narratives. Everything else is structured code with identical, auditable results.

Zero-AI Agents Exist

Several agents are entirely scripts and API calls. Data aggregation, compliance scheduling, portfolio monitoring triggers: pure automation with zero AI costs and zero hallucination risk.

Predictable, Auditable, Defensible

Compliance checks, advisor data routing, filing deadlines, and alert routing: identical output given identical input. No randomness, no hallucination, no drift. Audit-ready by design.

Most operations across the platform are deterministic automation; AI is used where judgment is required. This matters for fiduciary duty and regulatory defensibility.

The Human Approval Framework

Every advisor-facing and externally visible action requires human approval before reaching anyone outside your firm.

Alpheous never sends, publishes, or communicates externally without a human in the loop. For asset managers, this is not optional; it is a fiduciary requirement. The architecture enforces it.

Dual-Gate Approval

Advisor communications pass through the Compliance Review agent for SEC Marketing Rule validation, then to a human reviewer. Both gates must approve before anything sends.

Correction Learning

Edits are stored and used to improve future drafts. The system learns from your judgment, adapts to your firm's voice, but never bypasses the approval requirement.

Zero Unauthorized Communications

Advisor emails, meeting follow-ups, outreach drafts, marketing content, and LinkedIn posts sent without approval: zero. No code path bypasses this for external content. Work-queue follow-ups land in the wholesaler's own drafts folder; a draft that fails compliance is never written, and the log records who was skipped and why.

Every Surface Covered

Advisor emails, quarterly letters, market commentary, LinkedIn posts, website content, territory reports, and coverage dashboards all require compliance review plus human approval.

Approve From Anywhere, Recorded Everywhere

The approve, edit, or reject decision can be made in the Alpheous app or over Slack, email, or text. The channel changes; the rule does not. Nothing sends without that decision, and every decision is written to the audit log.

Trigger

Market data update, quarter end, advisor inquiry, or scheduled publication

Process

Pipeline gathers data, enriches context, routes to appropriate agent

Generate

AI drafts content with full context, market data, and firm voice

Compliance Gate

The Compliance Review agent validates SEC Marketing Rule compliance, disclaimers, performance claims

Human Review

Draft appears with full content, recipient, and compliance status, for approval in the Alpheous app or over Slack, email, or text

Action & Audit

Only on explicit approval does anything send. Full audit trail recorded.

3

Compliance—First Governance

Every agent operates from the same source of truth with SEC-aware governance enforced at the infrastructure level.

The biggest risk in multi-agent systems for asset managers is not just inconsistency; it is regulatory exposure. Alpheous solves this with unified compliance governance across all agents.

Unified Compliance Layer

Every agent operates under the same compliance framework: SEC Marketing Rule awareness, disclaimer requirements, performance claim validation rules, and recordkeeping obligations. The Compliance Review agent enforces these standards across every external output from every agent.

Policies Defined Once, Enforced Everywhere

Compliance constraints are defined once and enforced across all agents. 'All performance claims must reference source data and time periods' applies to the Research Writer, LinkedIn Agent, and every other agent that produces external content.

Preventing Regulatory Drift

Without governance, agents drift in tone, compliance rigor, and disclaimer usage. The governance layer keeps SEC compliance standards constant regardless of which agent produces the content or how many are operating simultaneously.

Advisor Data Isolation as Governance

Per-firm data boundaries enforced at the database level. Advisor records, CRM data, and communications are separated by infrastructure controls, not prompt instructions an AI might ignore. This isolation is a governance requirement, not just a security feature.

4

Agent Boundaries and Fiduciary Permissions

Every agent operates within defined permissions aligned with fiduciary duty. No agent can expand its own access.

Read/Write Separation

Each agent has explicit access boundaries designed for fiduciary compliance. The Research Writer accesses market data but cannot send advisor communications. The Compliance Agent reviews content but cannot alter fund performance data.

No Self-Expansion

Agents cannot expand their own permissions. New capabilities require supervised deployment with compliance review. Access is granted by humans through configuration, never requested at runtime.

Internal vs. External-Facing

Internal agents (research, data gathering, monitoring) produce no external output and need no approval. Advisor-facing agents require both compliance gate and human approval on every output.

SEC Recordkeeping Enforcement

Every agent action that touches advisor data or produces external content is logged with the detail required for SEC recordkeeping: which agent, which data, when, what output, who approved.

Cryptographic Accountability

Every action taken by every agent is recorded in hash-chained audit logs. Each entry carries a SHA-256 hash linking it to the previous entry, and each day's record is closed with a cryptographically signed root hash. For any regulatory inquiry, you walk the chain backward from the result to the trigger: the data that informed it, the compliance checks performed, the human who approved it. The chain is tamper-evident. Entries cannot be edited, inserted, or deleted after the fact without detection.

5

Trust Is Earned, Not Assumed

Autonomy increases only after demonstrated reliability through four phases. External-facing approval never goes away.

Shadow Mode

Week. Weeks 1–3

Agents observe and report but take no action. Every classification, recommendation, and draft is logged for review. Your team validates accuracy against your standards. This is how we deploy: maximum guardrails first, autonomy earned over time.

Supervised Execution

Week. Weeks 3–6

Agents generate outputs with 100% human approval required. The Compliance Review agent compliance gate is active on all external content. Corrections feed back into the system.

Calibrated Trust

Week. Weeks 6–10

Low-risk internal operations run with reduced oversight: market data gathering, research compilation, internal analysis. All advisor-facing actions still require compliance gate plus human approval.

Production Steady State

Week. Week 10+

Full capacity with permanent dual-gate approval on all advisor-facing and external communications. Internal research and data operations run autonomously. Monthly compliance posture reports document ongoing adherence.

External-facing approval never goes away. Every advisor email, outreach draft, market commentary, and LinkedIn post always requires compliance review plus human approval. Trust escalation applies only to internal operations. This is a fiduciary requirement, not a preference.

6

Structured in Practice

How Structured Agents blend deterministic code with targeted AI for asset management operations.

Market Commentary Pipeline

Agent. Investment Research Writer

STRUCTURED

- FOMC statement ingestion and parsing
- FRED economic data API polling
- Earnings data aggregation from verified sources
- Credit spread and yield curve monitoring
- Source attribution and citation tracking
- Research RAG pipeline for institutional sources

AI

- Market commentary drafting in firm voice
- Quarterly narrative synthesis from multiple data sources

SEC Marketing Rule Pipeline

Agent. Compliance Review

STRUCTURED

- Disclaimer library matching and insertion
- Performance claim detection and source verification
- Testimonial and endorsement classification
- Fair balance validation rules
- Recordkeeping metadata generation
- Compliance status tracking across all content

AI

- Nuanced compliance judgment on edge cases
- Context-aware disclaimer recommendations

Advisor Outreach Pipeline

Agent. Advisor Communications

STRUCTURED

- Advisor contact database management
- Territory coverage tracking
- Communication frequency tracking
- Outreach sequence scheduling
- Follow-up delivery confirmation logging
- CRM timeline update triggers

7

Distribution-Specific Governance

Territory boundaries, compensation isolation, and cross-wholesaler privacy enforced at the infrastructure level.

Territory Access Controls

Each wholesaler's view is scoped to their assigned territory. Advisor records, pipeline data, and outreach history outside their territory are architecturally inaccessible. Territory assignments are managed through configuration, not prompt instructions.

Compensation Data Isolation

Compensation structures, incentive calculations, and payout data are isolated from general agent queries. No agent can surface compensation information in messages, drafts, or reports without explicit configuration by a firm administrator.

Cross-Wholesaler Privacy

One wholesaler cannot see another's pipeline, outreach drafts, coverage statistics, or meeting notes by default. Cross-wholesaler visibility is available only to designated leadership roles and is logged in the audit trail.

Morning Brief Safeguards

Alph writes only the opening lines of each wholesaler's daily brief; everything else is the firm's own data or a linked source, and market headlines are linked to the publisher, never rewritten. Any sentence with a figure the brief cannot verify is removed before it sends. Every delivery is recorded in the compliance audit log.