February
2026

Know Your Agent (KYA)
A Framework for Trust and Accountability in the Agentic Enterprise

A White Paper by Aigentsphere

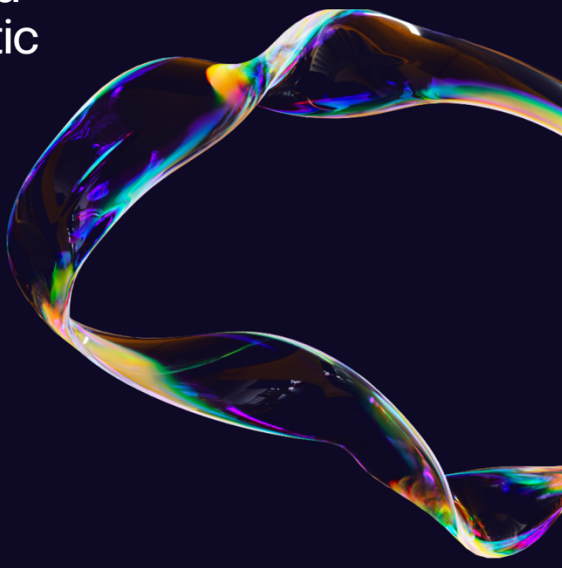

Aigentsphere

CONTACT US
info@aigentsphere.com
www.aigentsphere.com

Executive Summary

The enterprise is on the cusp of a new technological paradigm: the agentic era. A universe of AI agents is rapidly emerging, promising to automate complex workflows, drive unprecedented efficiency, and unlock new sources of value. However, this proliferation of autonomous, non-human actors creates a new and urgent governance challenge. Just as the financial industry developed "Know Your Customer" (KYC) frameworks to manage customer-related risks, enterprises must now adopt a **"Know Your Agent" (KYA)** framework to manage agent-related risks.

This white paper introduces KYA as a critical governance framework for the agentic enterprise. It draws a direct parallel to the principles of KYC, arguing that enterprises must be able to identify, verify, and continuously monitor every AI agent operating within their ecosystem. Without a robust KYA framework, organizations are exposed to a new and dangerous class of risks, including "shadow" AI, algorithmic bias, data breaches, and a complete lack of accountability for agent actions.

We propose a KYA framework built on three core pillars:

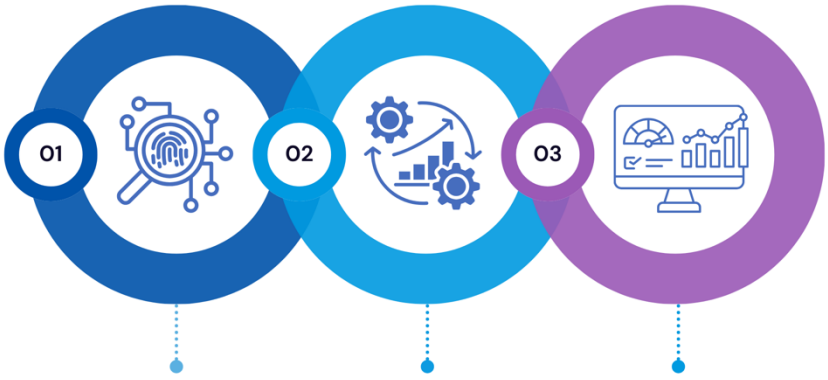

Agent Identification Program (AIP)
Establishing a definitive record of every agent's identity, provenance, and ownership.

Agent Due Diligence (ADD)
Assessing and classifying the risk profile of each agent before and during its lifecycle.

Continuous Agent Monitoring (CAM)
Maintaining ongoing visibility and control over agent behavior, performance, and cost.

Strategic Recommendation: The time to implement a KYA framework is now, before the proliferation of ungoverned agents creates an unmanageable and costly crisis. We strongly recommend the adoption of an independent, vendor-agnostic governance platform like **Aigentsphere**, which is purpose-built to deliver the comprehensive capabilities required for a robust KYA framework. Aigentsphere provides the central nervous system for KYA, enabling organizations to embrace the agentic era with confidence, trust, and control.

Table of Contents

Section 1: The Dawn of the Agentic Era and the Governance Imperative

The nature of work is changing. For decades, software has been a tool that humans use to perform tasks. Now, we are entering the agentic era, where software itself becomes an autonomous actor, capable of performing complex, multi-step tasks with minimal human intervention. These AI agents are not just chatbots or simple automation scripts; they are sophisticated digital workers that can reason, plan, and execute actions across multiple systems.

This represents a monumental opportunity for the enterprise. A recent report by Gartner predicts that by 2028, over 50% of large enterprises will have deployed AI agents in customer-facing and operational roles. McKinsey estimates that AI agents could unlock $4.4 trillion in annual economic value across industries. However, this rapid adoption is outpacing the evolution of enterprise governance. The core problem is one of identity and accountability. When an AI agent makes a mistake, who is responsible? When it accesses sensitive data, is it authorized? When it incurs costs, who is accountable?

Consider the practical reality facing enterprises today. A single organization might have dozens, hundreds, or even thousands of AI agents operating across different departments, platforms, and use cases. Some are customer-facing chatbots, others are internal productivity agents, and still others are sophisticated decision-making systems that approve loans, screen job candidates, or optimize supply chains. Each of these agents has the potential to create value, but also the potential to create risk. Without a comprehensive view of this agent ecosystem, organizations lack the fundamental visibility required for effective governance.

Without a framework to answer these questions, organizations are flying blind. This is where the concept of "Know Your Agent" (KYA) becomes essential. Borrowing from the battle-tested principles of "Know Your Customer" (KYC), KYA provides a structured approach to managing the risks associated with a non-human workforce.

> ***KYC vs. KYA: A Necessary Parallel***
>
> *KYC was created to prevent fraud, money laundering, and other financial crimes by verifying customer identities and monitoring their behavior. In the agentic enterprise, KYA is necessary to prevent algorithmic bias, data breaches, financial waste, and catastrophic errors by verifying agent identities and monitoring their behavior.*

The KYC-to-KYA Translation

A Framework Built on Proven Principles

The power of the KYA framework lies in its foundation on decades of proven KYC practices. The table below illustrates the direct parallels between these two governance frameworks:

	Financial Services Context		Agentic Enterprise Context	
01	Identity Verification	Confirm customer identity through documents, biometrics, and third-party validation	Agent Identification	Establish unique agent identifiers, track provenance, and document ownership
02	Risk Assessment	Classify customers by risk level based on transaction patterns and profile	Agent Risk Classification	Categorize agents by potential impact (low, medium, high, critical) based on capabilities and data access
03	Customer Due Diligence (CDD)	Understand the nature and purpose of customer relationships	Agent Due Diligence (ADD)	Assess agent purpose, capabilities, limitations, and potential business impact
04	Enhanced Due Diligence (EDD)	Apply stricter scrutiny to high-risk customers (e.g., politically exposed persons)	Enhanced Agent Due Diligence (EADD)	Apply stricter controls to high-risk agents (e.g., those making autonomous decisions affecting customers)
05	Ongoing Monitoring	Continuously monitor transactions for suspicious activity	Continuous Agent Monitoring (CAM)	Track agent behavior, performance, and costs in real-time to detect anomalies
06	Senior Management Accountability	KYC programs must be approved in writing by senior executives	Executive Governance	AI governance frameworks must have executive sponsorship and accountability
07	Audit Trails	Maintain comprehensive records of all customer interactions for regulatory compliance	Agent Audit Trails	Log all agent actions, decisions, and changes for compliance and forensic analysis
08	Periodic Review	Regularly update customer risk profiles based on new information	Dynamic Agent Profiling	Continuously re-assess agent risk profiles as behavior and context evolve

This parallel is not merely conceptual; it is operational. Just as financial institutions cannot operate without a robust KYC framework, enterprises in the agentic era cannot operate safely or compliantly without a robust KYA framework.

Section 2: The Three Pillars of a Know Your Agent (KYA) Framework

A robust KYA framework is built on three interconnected pillars, each designed to provide a specific layer of visibility and control over the AI agent lifecycle. These pillars work in concert to create a comprehensive governance structure that is both proactive and adaptive.

Pillar 1: Agent Identification Program (AIP)

The foundational pillar of KYA is the **Agent Identification Program (AIP)**. You cannot govern what you cannot see.

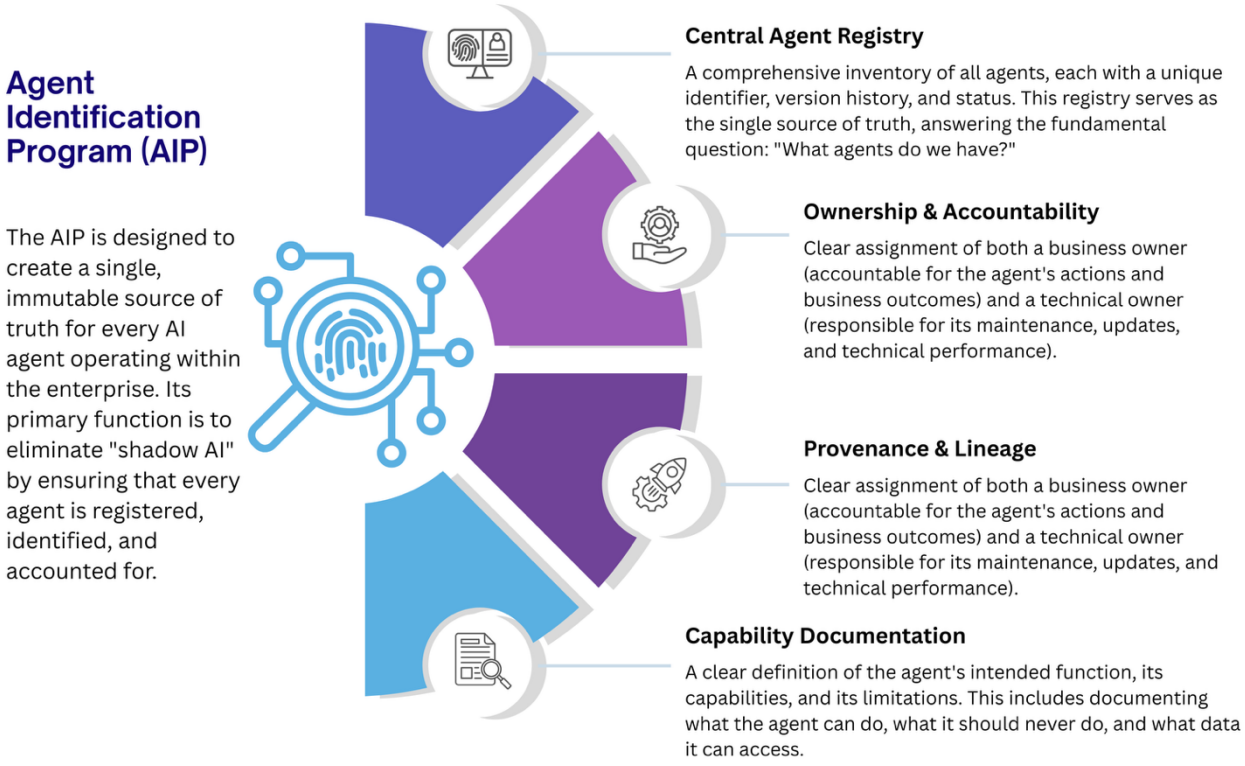

Real-World Scenario

A large financial services firm discovers through an internal audit that it has over 200 AI agents deployed across different business units, but only 60 are documented in any central system. The remaining 140 are "shadow agents"—undocumented, unmonitored, and ungoverned. Without an AIP, the firm has no visibility into what these agents are doing, who owns them, or what risks they pose. Implementing an AIP as the foundation of their KYA framework allows them to bring all agents into the light, establishing accountability and enabling proper governance.

Pillar 2: Agent Due Diligence (ADD)

Once an agent's identity is established, the next step is to understand its risk profile. **Agent Due Diligence (ADD)** is the process of assessing and classifying the potential risks associated with each agent before it is deployed. This is not a one-time check, but an ongoing process that adapts to changes in the agent's behavior and the business context.

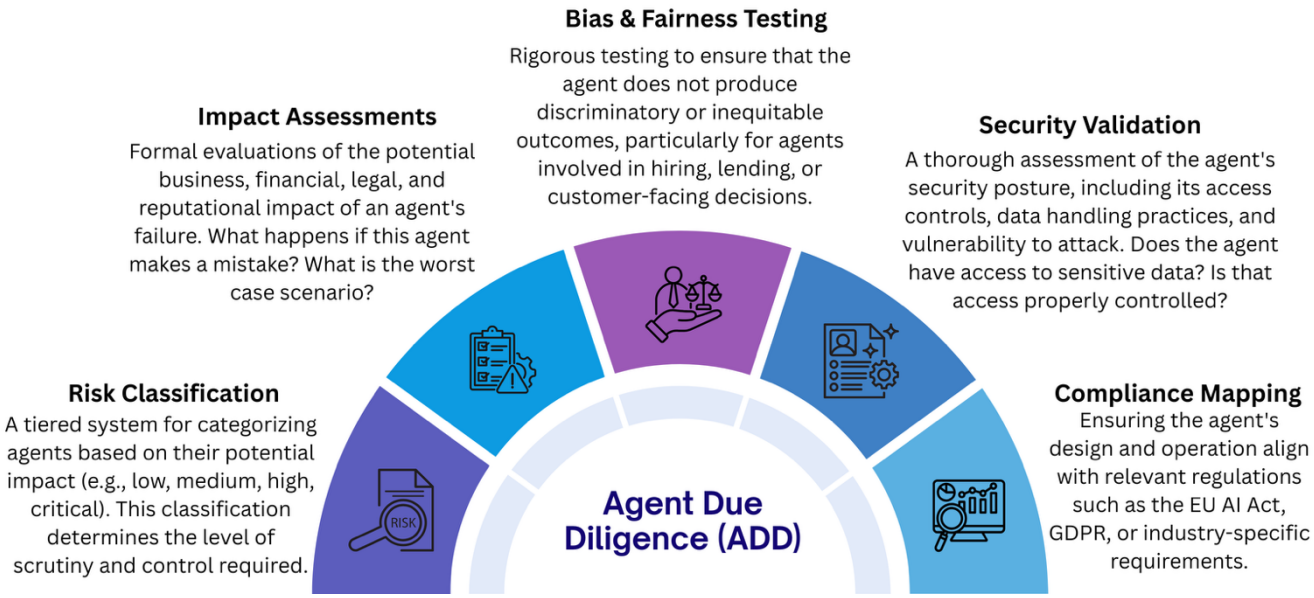

For high-risk agents, **Enhanced Agent Due Diligence (EADD)** is required, mandating stricter controls, human-in-the-loop approval workflows, and more frequent reviews. Examples of high-risk agents include those that make autonomous decisions affecting customers (loan approvals, job candidate screening), those that process sensitive personal data (health records, financial information), or those that interact directly with the public on behalf of the organization.

Real-World Scenario

A healthcare organization deploys an AI agent to triage patient inquiries and recommend whether they should seek immediate care. Under a KYA framework, this agent would be classified as "critical risk" due to its potential impact on patient safety. EADD would require bias testing to ensure equitable treatment across demographic groups, security validation to protect patient data, and human-in-the-loop workflows to ensure a licensed healthcare professional reviews high-risk recommendations before they are communicated to patients.

Pillar 3: Continuous Agent Monitoring (CAM)

The final pillar of KYA is **Continuous Agent Monitoring (CAM)**. AI agents are not static; they are dynamic and can evolve over time. CAM provides the real-time visibility and control necessary to manage agents in production, ensuring they operate within their intended boundaries and continue to deliver value.

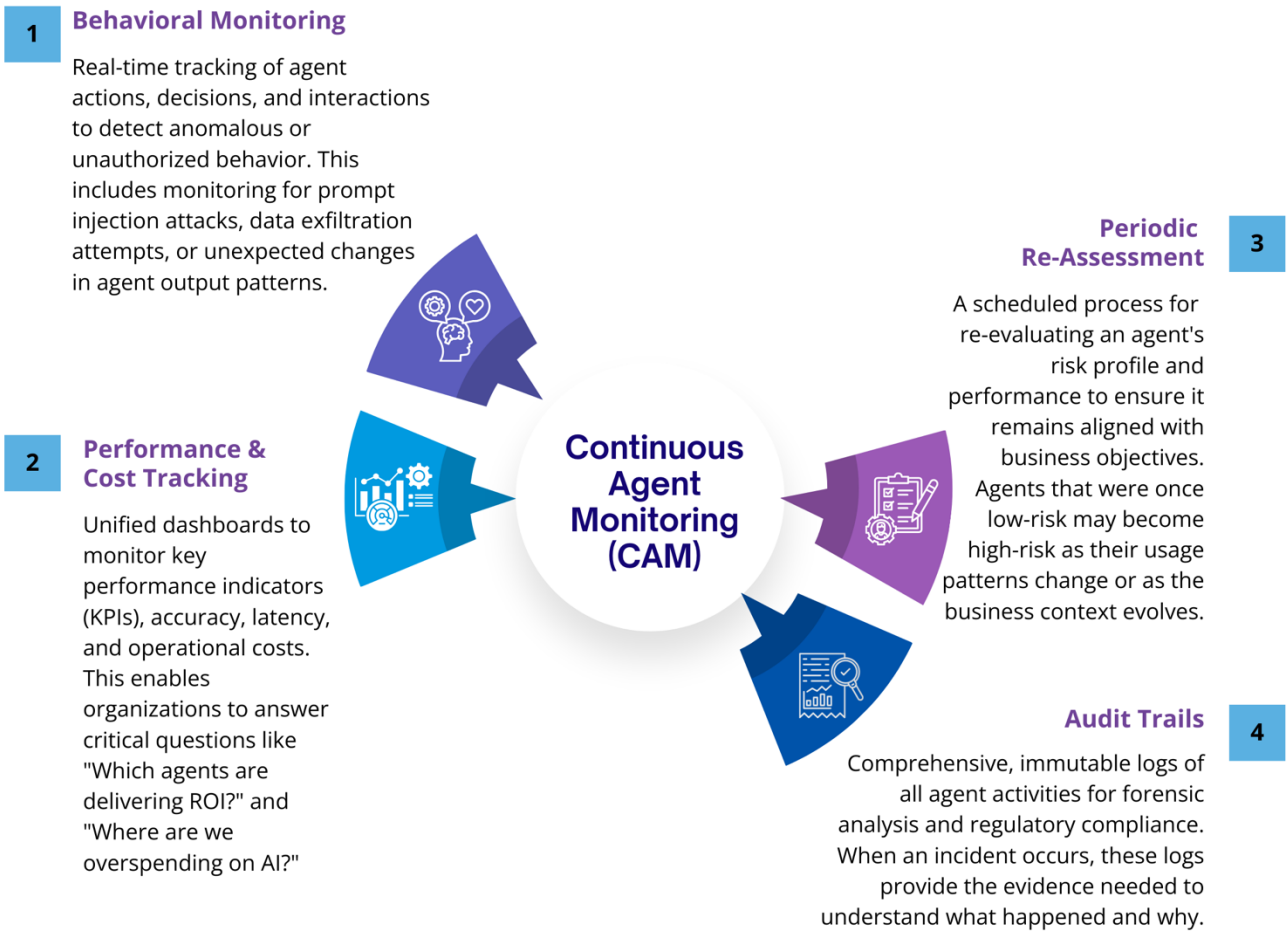

Real-World Scenario

An e-commerce company deploys an AI agent to handle customer service inquiries. Initially, the agent performs well, but after three months, the monitoring system detects a significant increase in customer complaints and a spike in API costs. Investigation reveals that a recent update to the underlying language model has caused the agent to provide longer, more verbose responses that are less helpful and more expensive. Without CAM, this degradation would have gone unnoticed for months, eroding customer satisfaction and wasting budget. With CAM, the issue is identified within days, and the agent is rolled back to the previous version while a fix is developed.

Section 3: Aigentsphere: The Purpose-Built Platform for KYA

The Know Your Agent framework requires a new class of enterprise software—one that is independent, vendor-agnostic, and purpose-built for governance. **Aigentsphere** is this platform. It provides the comprehensive Layer 3 (Central Governance) and Layer 4 (Risk Management) capabilities that are essential for implementing a robust KYA framework.

Here is how Aigentsphere's core modules map directly to the three pillars of KYA:

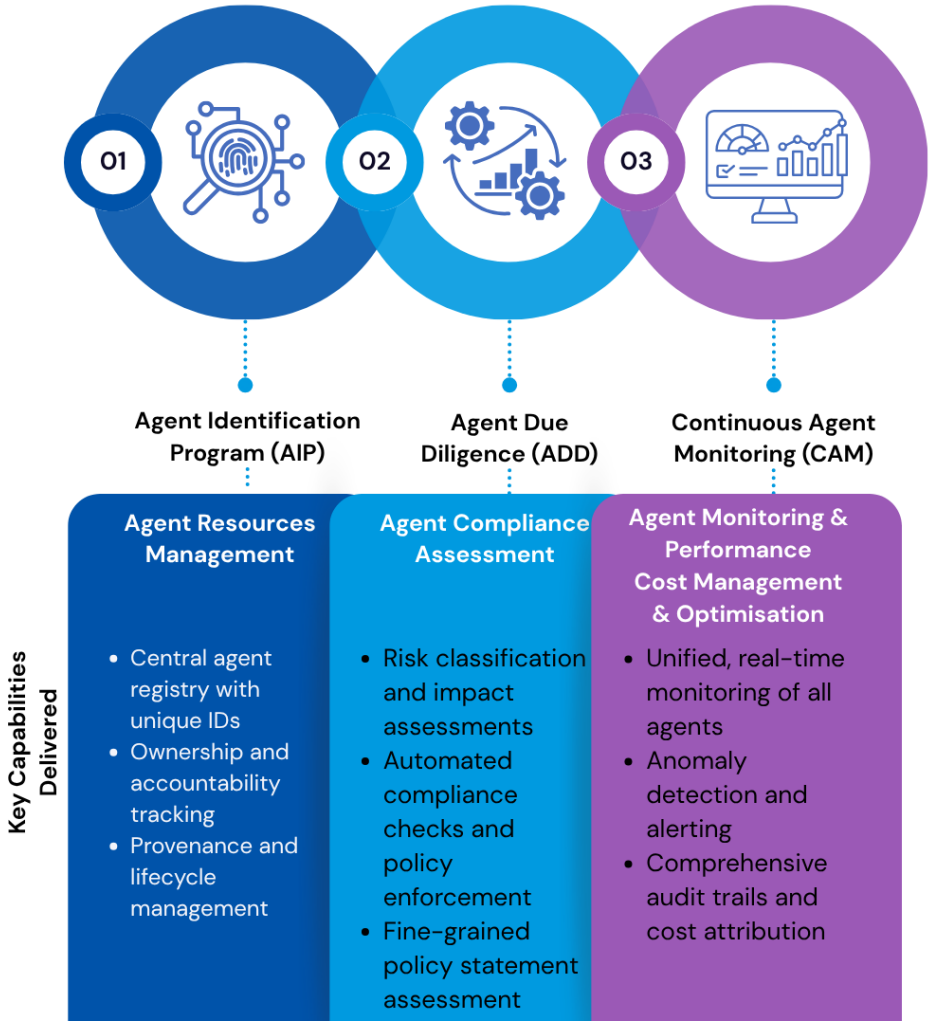

Why Aigentsphere is the Essential KYA Platform:

- **Independence:** As a vendor-agnostic platform, Aigentsphere provides objective, unbiased oversight of all AI agents, regardless of the underlying development platform (OpenAI, Google, Anthropic, etc.).
- **Completeness:** Aigentsphere is the only platform that provides a comprehensive, integrated solution for all three pillars of KYA.
- **Scalability:** Aigentsphere is designed to scale from a handful of agents in a pilot project to an enterprise-wide ecosystem of thousands of digital workers.
- **Auditability:** With its comprehensive logging and reporting capabilities, Aigentsphere provides the audit-ready documentation required to demonstrate compliance with emerging AI regulations.

Section 4: Implementing Your KYA Framework: A Phased Approach

Implementing a comprehensive KYA framework is not an all-or-nothing proposition. Organizations can adopt a phased approach that delivers immediate value while building toward full maturity. Aigentsphere is designed to support organizations at every stage of this journey.

Phase 1: Discovery & Baseline (Weeks 1-4)

Objective: Establish visibility into the current agent landscape and create a baseline inventory.

Key Activities:

- Deploy Aigentsphere's Agent Resources Management module
- Conduct an enterprise-wide agent discovery process
- Register all identified agents in the central registry
- Assign initial ownership and accountability
- Document known capabilities and limitations

Deliverable: A comprehensive inventory of all AI agents, eliminating shadow AI and establishing the foundation for governance.

Phase 2: Risk Assessment & Classification (Weeks 5-8)

Objective: Assess and classify the risk profile of each agent to prioritize governance efforts.

Key Activities:

- Develop risk classification criteria aligned with business objectives
- Conduct impact assessments for all registered agents
- Classify agents into risk tiers (low, medium, high, critical)
- Identify agents requiring Enhanced Agent Due Diligence (EADD)
- Deploy Aigentsphere's Training & Compliance and Access Controls modules

Deliverable: A risk-stratified agent portfolio with clear governance requirements for each tier.

Phase 3: Monitoring & Control (Weeks 9-12)

Objective: Implement real-time monitoring and establish ongoing governance processes.

Key Activities:

- Deploy Aigentsphere's Monitoring & Performance and Cost Management modules
- Configure dashboards and alerting for high-risk agents
- Establish baseline performance and cost metrics
- Implement automated compliance checks and policy enforcement
- Train business and risk teams on using the platform

Deliverable: A fully operational KYA framework with real-time visibility and control over all agents.

Phase 4: Optimization & Maturity (Ongoing)

Objective: Continuously refine and optimize the KYA framework based on learnings and evolving business needs.

Key Activities:

- Conduct periodic re-assessments of agent risk profiles
- Optimize agent performance and cost efficiency
- Expand governance to new agents and use cases
- Leverage the Agent Marketplace to promote reuse of vetted agents
- Integrate KYA processes into standard development workflows

Deliverable: A mature, embedded governance culture where KYA is the default operating model for all AI initiatives.

The agentic era is not a distant future; it is happening now.

Enterprises that fail to implement a robust governance framework will inevitably face a crisis of control, compliance, and cost. The "Know Your Agent" framework provides a clear and actionable path forward, enabling organizations to embrace the power of AI agents while mitigating the inherent risks.

Retrofitting governance onto a chaotic and ungoverned agent ecosystem is a recipe for failure.

The time to act is now. By implementing a KYA framework and a purpose-built governance platform like Aigentsphere from day one, organisations can build a foundation of trust, accountability, and control that will enable them to thrive as an agentic enterprise.

About Aigentsphere

Aigentsphere is the leading independent AI governance platform, purpose-built to provide comprehensive Layer 3 (Central Governance & Performance Management) and Layer 4 (Risk Management & Human Oversight) capabilities for enterprise AI. Our vendor-agnostic platform enables organizations to implement the Know Your Agent (KYA) framework from day one, accelerating innovation while maintaining control, compliance, and trust.

Founded on the principle that effective AI governance requires independence from underlying AI development platforms, Aigentsphere provides a single, unified view and control point for all AI agents across the enterprise, delivering the complete capabilities required for a robust KYA framework.

Aigentsphere is designed to scale from pilot projects to enterprise-wide deployments, supporting organizations at every stage of their AI maturity journey. With comprehensive audit trails, real-time monitoring, and automated compliance checks, Aigentsphere provides the foundation for responsible, scalable, and trusted AI agent deployment.

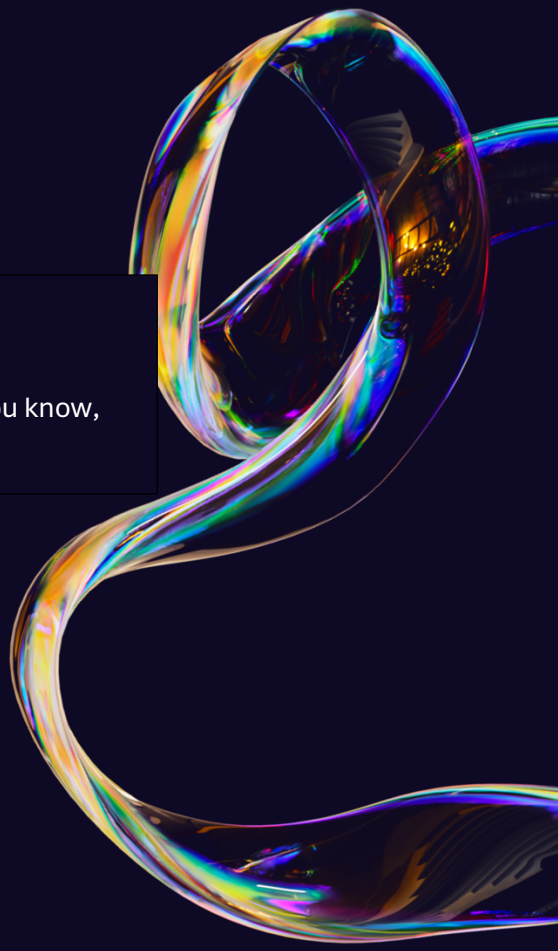

Ready to implement your KYA framework?

Visit www.aigentsphere.com or contact our team (info@aigentsphere.com) to learn how Aigentsphere can help you know, govern, and optimize every AI agent in your enterprise.