CEREVANT INTELLIGENCE / RESEARCH 002 / OCTOBER 2026

Beyond Know Your Agent

The Enterprise Governance Architecture for Autonomous AI



Identity. Authority. Context. Accountability.

Executive Perspective

The enterprise AI conversation is entering a new phase. The central question is no longer simply whether artificial intelligence can perform complex tasks. It is whether organizations can establish sufficient trust, control, and accountability to permit autonomous systems to execute those tasks on their behalf.

In his October 22, 2026 Financial Times article, *"'Know Your Agent' will be the new rule of the AI age,"* John Thornhill introduces Know Your Agent (KYA) as an emerging governance imperative, drawing a parallel with the established Know Your Customer (KYC) obligations of financial institutions.

The article identifies three fundamental challenges: establishing agent identity, controlling agent behaviour, and determining accountability when autonomous actions produce unintended consequences.

Cerevant Intelligence extends this discussion into the enterprise operating environment.

Our central proposition: Agent identity establishes who is acting. Enterprise governance must establish whether that action is authorized, contextually appropriate, observable, and accountable.

01 | From Digital Assistants to Autonomous Actors

Traditional enterprise applications operate within relatively predictable workflows. Agentic systems introduce a different operating model: software capable of interpreting objectives, selecting tools, accessing systems, and executing sequences of actions with varying degrees of independence.

This changes the nature of enterprise risk.

An agent may execute an instruction correctly at the technical level while producing an unacceptable commercial, operational, or regulatory outcome.

The challenge is therefore not merely securing access to systems. It is governing decisions made within those systems.

02 | The Five Layers of Agentic Trust

Cerevant proposes a five-layer conceptual architecture for evaluating enterprise agent deployments.

Identity: Establish verifiable agent identity, provenance, ownership, and lifecycle management.

Authority: Define permissions, transaction limits, delegated decision rights, and escalation requirements.

Context: Incorporate business rules, organizational dependencies, exceptions, and operational consequences into agent decision-making.

Observability: Maintain traceable records of agent actions, tool usage, policy decisions, exceptions, and human interventions.

Accountability: Establish responsibility for agent deployment, supervision, incidents, remediation, and outcomes.

These layers are interdependent. Strong authentication without contextual safeguards cannot guarantee appropriate behaviour. Likewise, detailed activity logs cannot substitute for clearly assigned operational responsibility.



Cerevant conceptual model. Interdependent controls with dependencies and feedback. Not an FT framework or a regulatory standard.

03 | Context: The Missing Enterprise Control Layer

Consider an AI agent tasked with improving working-capital efficiency across a hospitality portfolio.

The agent may identify overdue invoices, renegotiate payment schedules, and recommend delayed expenditures.

From a narrowly defined financial perspective, these actions could appear successful.

Yet an agent unaware of contractual obligations, critical maintenance dependencies, franchise requirements, or vendor relationships may optimize one metric while creating disproportionate operational risk elsewhere.

This illustrates a central challenge of autonomous enterprise systems: **local optimization does not necessarily produce organizational optimization.**

Cerevant identifies contextual intelligence as a critical bridge between technical automation and responsible enterprise decision-making.

04 | Governance Must Become Operational

Enterprise agent governance cannot remain exclusively a policy document.

It must become part of the technical and operational architecture through which agents are deployed, supervised, and evaluated.

This requires explicit control boundaries, graduated autonomy, human approval for consequential decisions, continuous evaluation, incident response procedures, and mechanisms for suspending or revoking agent privileges.

Organizations should also distinguish between agents that recommend actions, agents that execute reversible actions, and agents capable of initiating consequential or irreversible changes.

Each category requires a different governance threshold.

05 | From Agent Adoption to Agent Assurance

As enterprises scale autonomous systems, adoption metrics alone become insufficient.

Agentic maturity should increasingly be assessed through evidence of control effectiveness.

Relevant measures include the proportion of agents with assigned owners, completeness of authorization policies, frequency of unauthorized-action attempts, escalation effectiveness, traceability of consequential decisions, and time required to contain an incident.

These measures do not eliminate autonomous-system risk. They make it more visible, measurable, and manageable.

06 | The Cerevant Perspective

The next competitive advantage in enterprise AI may not belong solely to organizations deploying the greatest number of agents.

It may belong to those capable of coordinating autonomous intelligence with institutional knowledge, defined authority, reliable infrastructure, and demonstrable accountability.

The transition from AI experimentation to operational autonomy demands more than model capability.

It demands an architecture of trust.

Knowing your agent is the beginning. Governing its authority, understanding its context, and accounting for its actions is the enterprise imperative.

Source Acknowledgment

Primary editorial inspiration: John Thornhill, *"'Know Your Agent' will be the new rule of the AI age,"* Financial Times, October 22, 2026.

This publication is an independent Cerevant Intelligence analysis extending the article's governance themes into a proposed enterprise operating framework. The Five Layers of Agentic Trust are presented as Cerevant's conceptual model, not as an FT framework or established regulatory standard.

CEREVANT INTELLIGENCE

Research & Advisory | Enterprise AI · Autonomous Systems · Digital Governance