

THE AI INSTITUTE / OPERATING INTELLIGENCE

MONDAY BRIEF · 14 SEP / 2026



AI Agents Need One Enterprise Control Layer

Vendors are converging on shared infrastructure for context, authority, oversight and cost. Boards should decide who owns it before agent sprawl does.

Boards, C-suite and senior management

Research period: September 2026 – September 2026

First published: 14 September 2026

Current to 14 September 2026

PUBLICATION RECORD

An Institute thesis and decision framework

The market is converging on an agent control layer; boards should decide which functions are shared and who owns them before individual product teams and suppliers define the operating model by default.

WHAT LEADERS SHOULD TAKE FROM THIS PAPER

- 01

Treat identity, authority, context, evaluation, monitoring, cost and exit as shared enterprise infrastructure.
- 02

Assign one executive owner for the control layer without concentrating every operating decision in one team.
- 03

Fund common controls before the number of agent products and workflows multiplies.
- 04

Require every supplier to connect to the enterprise record and preserve export, audit and withdrawal paths.
- 05

Adapt the architecture to local law, infrastructure and institutional capacity while keeping the decision standard consistent.

Research method

This Monday Brief synthesises material developments published from 8 to 14 September 2026, led by official announcements from agent, enterprise-software, integration and network suppliers and checked against NIST standards work, a bounded arXiv deployment study, current infrastructure and delivery investment, public-sector programmes and the preceding Institute catalogue. The public treatment translates the technical convergence into an enterprise ownership and capital decision.

First published 14 September 2026 · Updated 24 September 2026 · Research period September 2026 – September 2026 · Version 1.1 · Suggested citation: The AI Institute, *AI Agents Need One Enterprise Control Layer* (2026).

CONTENTS

Inside this paper

01	The product around the agent is becoming the strategic asset The Board Signal
02	The market is converging on a common enterprise problem What Changed
03	Make seven functions common before agents multiply Shared Infrastructure
04	Give the layer one owner without creating one gatekeeper Accountability
05	Fund reuse, but price the concentration risk Capital Allocation
06	Measure the human-agent system, not the model alone Operating Performance
07	Keep the decision standard global and the architecture local Global Reality
08	Make the control layer visible before approving more agents The Next 30 Days
M	Method and limitations
R	References

How to use this brief

Start with the decision and use the framework to challenge the current plan. Figures retain their original populations and definitions. Institute analysis is labelled, and limitations are recorded before the references.

The product around the agent is becoming the strategic asset

The most important AI development of the week is not a new model. It is the sudden visibility of the layer that sits around the model and lets an agent continue working: the context it can use, the tools it can reach, the actions it may take, the record it leaves and the conditions that make it stop.

OpenAI now offers its agent harness through an API. Salesforce has announced an enterprise harness and control plane. Boomi is selling central control across agents and core systems. Zayo has opened a governed route for approved agents to act on network operations. Different suppliers use different language, but they are competing for the same position: the layer through which enterprise AI understands, acts and is managed. ¹²³⁴

This is a board decision because the layer can become more durable than any individual model. It will accumulate policy, workflow history, permissions, evaluation results and cost information. If it is designed one product at a time, the organisation will inherit an operating model assembled by purchasing decisions. If it is designed as shared infrastructure, leaders can decide what remains common as models and suppliers change.

The durable AI asset may be the control layer, not the model.

The market is converging on a common enterprise problem

On 10 September, OpenAI launched the Agents API in public beta. It hosts and maintains the harness while allowing the compute environment to sit in an OpenAI-managed sandbox, the customer's own infrastructure or a partner environment. The offer turns long-running context, tool use and coordination into a managed service rather than a component every team must assemble.¹

A day later, Salesforce described six capabilities across context, agency, action, governance, security and models, together with an AI Control Plane intended to register agents, apply identity and policy, evaluate performance, observe behaviour and outcomes, and control cost across Salesforce and third-party systems. Boomi has made a similar case from the integration layer, while Zayo has moved the idea into network operations, where approved agents can interpret conditions and take authorised action.²³⁴

These are supplier announcements, not a settled market standard. Availability, pricing, interoperability and independent customer outcomes remain incomplete. But the convergence is strategically useful. It says that organisations moving beyond isolated assistants will need common services around them. The decision is no longer whether that layer exists. It is whether the enterprise chooses it deliberately or acquires it by accident.

Make seven functions common before agents multiply

First, keep one register. The organisation should know which agents exist, which business process each serves, who owns it, which supplier and model it uses, and whether it is experimenting, operating or retired. An unregistered agent is not a productivity tool; it is an unknown operating dependency.

Second, give every agent its own identity and bounded authority. It should not borrow a person's standing access or repeatedly ask for approvals until someone clicks through. NIST warns that agentic systems are repeating familiar identity mistakes, including consent fatigue and routes that can encourage users to surrender sensitive information. Its broader review of stakeholder responses also found widespread agreement that familiar security practices need adaptation for agents. Existing identity discipline still applies, but continuous software action makes scope, duration and withdrawal more important.⁵¹¹

Third, separate business context from the model. The definitions, policies, customer commitments and workflow state that make an answer useful belong to the enterprise. They need a governed route into the agent and a portable record when the model changes.

Fourth, use one evaluation language. Each job can retain its own finish line, but reliability, human review, cost, incidents and unresolved exceptions should roll into a comparable enterprise view. Fifth, observe action rather than prompts alone. Leaders need to see what the agent attempted, which system changed, what failed and who accepted the outcome.

Sixth, manage total cost. Model usage is only one component; tool calls, infrastructure, human review, integration and rework belong beside it. Seventh, preserve exit. Policies, workflow definitions, evaluation cases, logs and business records must remain available when a supplier, model or control product changes.

EXHIBIT / THE SEVEN SHARED FUNCTIONS

1 Register Know every agent, job, owner and lifecycle state. Enterprise inventory	2 Identity Give software bounded, withdrawable authority. Security and access	3 Context Keep business meaning governed and portable. Data and knowledge	4 Evaluate Compare reliability, review and full operating cost. Assurance
5 Observe Record actions, changes, failures and acceptance. Operations	6 Cost Manage usage, infrastructure, review and rework. Finance	7 Exit Export the rules, records and workflow state. Procurement	

Give the layer one owner without creating one gatekeeper

The control layer needs an executive owner because fragmented ownership produces fragmented authority. The likely accountable leader is the CIO or COO, depending on whether the organisation treats agents primarily as technology infrastructure or as a new form of operating capacity. The CISO, data leader, risk function, procurement and workforce leadership should hold defined decision rights inside the design.

One owner does not mean one team approves every action. Central control should provide the register, identity pattern, minimum operating record, evaluation framework, cost view and withdrawal mechanism. Business functions should still define the outcome, permitted work, acceptable result and exceptions that require human judgement. This is a federation: common rails with local accountability.

The board's role is to approve the boundary. It should ask which functions must be common, which may vary by business unit and which decisions can never be delegated to a supplier. Without that boundary, a platform team can become an accidental policy-maker or a product team can grant authority that the enterprise has never consciously accepted.

Centralise the rails; keep the business decision with the accountable function.

EXHIBIT / A FEDERATED OWNERSHIP MODEL

BOARD	PLATFORM	BUSINESS
Boundary Approve what must be common and what cannot be delegated. Enterprise accountability	Rails Provide identity, records, evaluation and withdrawal. CIO or COO	Decision Own the outcome, permitted work and accepted result. Accountable function

Fund reuse, but price the concentration risk

A shared layer can remove real cost. Product teams should not each build a separate agent register, identity system, logging format, evaluation service and cost dashboard. Reuse shortens the route from a controlled experiment to an operating workflow and gives risk, security and finance one view across the portfolio.

The same layer can create deep dependency. A supplier that holds the context, policies, evaluation history and action records may become harder to replace than the model itself. The switching question therefore belongs in the original investment case, not the renewal negotiation. Buyers should require documented interfaces, export tests, independent logs and a credible customer-managed or alternative operating path for consequential work.

Current market investment shows why the choice will arrive quickly. Accenture and Google Cloud plan a dedicated Gemini Enterprise group with 1,000 forward-deployed engineers, while Oracle reported triple-digit cloud-infrastructure growth and more than US\$30 billion of additional AI cloud contracts in its latest quarter.⁷⁸ Delivery capacity and infrastructure demand are rising. They do not tell a board which control layer will create value or remain portable.

Use two budgets. The first funds the common layer as operating infrastructure. The second funds job-level deployments that must use it. This makes duplicated controls visible and prevents a successful pilot from silently carrying its local workaround into enterprise scale.

EXHIBIT / THREE CAPITAL QUESTIONS

REUSE	LOAD	EXIT
What becomes common? Fund the control once when many workflows need it. Platform budget	What does operation cost? Include human review, incidents, integration and rework. Job budget	What survives a supplier change? Test export and replacement before dependency grows. Portfolio risk

Measure the human-agent system, not the model alone

A model score can hide the operating burden around it. A recent research preprint offers a useful bounded example. Across 16 agent systems and 750 clinical-audit cases, two systems with almost the same autonomous accuracy required different levels of human review—39.2% and 29.6%—to reach the same stated reliability target under the evaluated policy.⁶

The percentages are not a global benchmark and the paper is not peer reviewed. The newest available arXiv batch also remains dominated by early work on coordination, memory, stability and evaluation rather than representative enterprise outcomes.¹² The management lesson is stronger than the number: similar-looking agents can create materially different supervision costs. A control layer should therefore record when people intervene, which cases are deferred, how long review takes and whether the accepted result improves the business process.

This changes the buying conversation. Ask suppliers to demonstrate the complete operating point for the defined job: reliability, permitted authority, human attention, time, total cost and failure recovery. If the answer stops at task completion, the organisation still does not know whether it can run the system responsibly.

Keep the decision standard global and the architecture local

The seven functions travel across markets, but their implementation does not. In the United States, supplier platforms and NIST's standards work may push the market toward reusable identity and interoperability patterns. Public agencies are also receiving new pricing and adoption support. In the European Union, public-service pilots begin inside a different mix of procurement, sovereignty, data protection and AI Act obligations.⁵⁹¹⁰

Regulated sectors add another layer. The United Kingdom's new healthcare-commission blueprint is a recommendation rather than enacted reform, but it reinforces the likelihood that assurance, monitoring and accountability will vary by sector. UNESCO's global forum can shape readiness language without creating national law. A multi-market control layer must record jurisdiction, system role and sector rather than treating policy as one universal setting.¹³¹⁴

In markets where connectivity, compute, specialist labour or local-language resources are constrained, the layer may be federated, partly manual or shared through public infrastructure. That does not lower the need for identity, bounded authority and an operating record. It changes the affordable way to provide them.

Organisations should connect the layer to existing privacy, cyber, public-sector assurance and critical-infrastructure obligations, and keep the same buyer-side ownership even when technology is hosted offshore. The common board question is simple: can the organisation still see, constrain and replace the agent under its actual local conditions?

Make the control layer visible before approving more agents

First, inventory the ten most consequential agent workflows, including pilots. Record the job, owner, supplier, model, data, tools, authority and current lifecycle state. Do not wait for a perfect enterprise catalogue; the first pass is meant to expose variation and unknowns.

Second, choose the seven common functions and assign their owners. Decide which existing identity, data, security, observability, finance and procurement systems can be extended rather than replaced. The aim is not a new technical empire. It is one coherent enterprise record and one way to withdraw authority.

Third, run one comparison. Take two agent workflows that appear similar and compare their reliability, human review, cost, incidents and export path. This will show whether the common layer is reducing operating friction or merely adding another platform.

Finally, change the next approval paper. Any material agent proposal should state how it joins the enterprise register, receives identity, obtains context, records action, proves acceptable operation, reports full cost and exits. If those answers remain product-specific and invisible to the institution, the proposal is not ready to scale.

The board question for this week is not which agent is most impressive. It is whether the organisation has decided what every agent must have in common before the estate grows beyond anyone's view.

Before approving another agent, decide what every agent must have in common.

RESEARCH RECORD

Method and limitations

Method

This Monday Brief synthesises material developments published from 8 to 14 September 2026, led by official announcements from agent, enterprise-software, integration and network suppliers and checked against NIST standards work, a bounded arXiv deployment study, current infrastructure and delivery investment, public-sector programmes and the preceding Institute catalogue. The public treatment translates the technical convergence into an enterprise ownership and capital decision.

Limitations

The selected product announcements are first-party supplier sources and do not establish typical customer outcomes, interoperability or full availability. The READY study is a preprint using one clinical-audit workflow and a researcher-defined reliability target. Oracle's result establishes supplier demand, not customer returns. Regional examples differ in law, sector, infrastructure, language and institutional capacity; no market is treated as a global adoption proxy.



Sources and limitations

Source facts link to the original publication. Institute analysis reconciles definitions and turns the material into a management proposition. Frameworks are decision aids, not claims of universal causality. Read every important statistic with its population, date, definition and caveat.

Research cutoff: 14 September 2026. Review cycle: six months, or earlier after a material change in the underlying facts, standards or policy.

REFERENCES

Evidence behind the thesis

01 OpenAI, Introducing the Agents API

Official product announcement dated 10 September 2026; public beta and selected customer cases.

02 Salesforce, Trusted Enterprise AI Harness

Official announcement dated 11 September 2026; availability, pricing and independent outcomes remain incomplete.

03 Boomi, Enterprise AI Control Plane

Official product announcement dated 2 September 2026; supplier claims require customer validation.

04 Zayo, Agentic Networking

Official announcement dated 8 September 2026; production claim is not an independent customer outcome.

05 NIST, Why Agentic AI Needs a Strong Identity Foundation

Standards-agency analysis dated 27 August 2026; guidance and project work, not enacted law.

06 arXiv, READY or Not: Reliable Enterprise Agent Deployment

New preprint submitted 2 September 2026; 16 systems and 750 cases in one clinical-audit workflow.

07 Accenture and Google Cloud, Gemini Enterprise Business Group

Joint corporate announcement dated 8 September 2026; planned workforce and selected customer case.

08 Oracle, Q1 FY2027 results

Issuer release dated 10 September 2026; supplier demand does not establish customer returns.

09 OpenAI and GSA, expanded US government access

Official supplier agreement dated 10 September 2026; one-country access offer and outcomes pending.

10 European Commission, GenAI pilots for public administration

Official EU programme event on 14 September 2026; pilots began 1 July and outcomes are pending.

11 NIST, Security Considerations for AI Agents

Official response analysis published 18 May 2026; stakeholder synthesis rather than a mandatory standard.

12 arXiv, recent cs.AI submissions

Newest available batch is 11 September 2026 with 171 entries; mixed status and mostly preprints.

13 UNESCO, Global Forum on the Ethics of AI

Official forum programme for 14–17 September 2026; voluntary normative framework.

14 UK Government, AI in healthcare regulatory blueprint

Official recommendation announcement dated 10 September 2026; not enacted reform.

THE AI INSTITUTE

Make the decision.
Then measure what changes.

Research · Readiness · Capability

theai.institute