



OCTOBER 2026

Agentic connectivity

Why integration matters more than ever



Introduction

Agents are the next enterprise channel, a new class of consumer that reads data, makes decisions, and acts on its own. But AI agents are only as capable as the systems they can reach. As enterprises move agents into production, the integration layer beneath them becomes the difference between agents that generate value and those that create risk, and many organizations are adding agents faster than they are adding the infrastructure to govern them.

Agents deployed as isolated projects can lead to agent sprawl, with duplicated connections, inconsistent security, and no central view of what agents are doing or why. Agents that act on stale, fragmented, or poorly governed data can turn autonomy into a liability, making confident decisions on the wrong information. And disconnected agents can add cost by rebuilding connectivity the enterprise has already solved elsewhere.

Organizations should treat a trusted data foundation as the starting point for enterprise-scale agents, with integration as the connective layer that makes that foundation usable. By bringing together Salesforce for business context, MuleSoft for governed connectivity, Informatica for trusted, high-quality data, and Deloitte Digital for transformation architecture and governance, organizations can give agents governed access to real-time data through a shared connectivity layer. This enables them to move from AI experimentation to production with the security, observability, and governance that boards and regulators now expect. This is the same closed loop principle behind Deloitte's broader point of view on enterprise architecture: insights only create value when they flow into governed action, and outcomes continuously refine the data and decisions that follow.

The integration question behind the AI agent wave

With autonomous AI agents becoming a standard feature in the enterprise technology stack, enterprise software is in the middle of one of its most significant architectural changes since the move to the cloud. Salesforce and other major platform vendors have introduced agent capabilities in the past year, and enterprise AI spending on agentic use cases is accelerating.

That acceleration has brought skepticism about whether organizations still need the middleware, API layers, and governance structures they've spent years building.

If agents can retrieve data from any system on demand, is traditional integration architecture still necessary?

This assumption overlooks how much agents depend on the systems around them. Agents need governed access to enterprise data, defined interfaces with proper permissions, and reliable communication layers to coordinate with one another. Those needs grow in an agentic environment, because each new agent adds connections to the enterprise architecture and increases the demands on integration infrastructure.



The connectivity challenge is about to get harder

[The 2026 Connectivity Benchmark Report by Deloitte](#)

[Digital & MuleSoft](#) found that the average enterprise now runs 957 applications, and that organizations furthest along in their agentic journey run 1,057 on average, 10% more than the overall mean. Yet only 27% of those applications are integrated. Managing the connections between those systems (through iPaaS platforms, API management, and data pipelines) is already a resource-intensive challenge for enterprise IT.

Now consider that nearly all of these applications are gaining their own AI agent, and that organizations are simultaneously building custom agents for internal use cases. In the Connectivity Benchmark, 88% of organizations have already adopted agentic transformation to some extent, and 98% plan to do so. Each agent needs to communicate with enterprise systems, data stores, and increasingly with other agents. An environment with roughly a thousand applications and a comparable population of agents creates a connectivity surface that is exponentially larger and more complex than what most integration teams are managing today.

Without deliberate architecture, organizations may end up with the same tangle of point-to-point connections that plagued enterprise IT before microservices and modern integration platforms brought order to the problem, a failure mode the market has started calling “agent sprawl”.





Why agents fail without integration and real-time data

Agents are autonomous software components that gather information, make decisions, and take actions. Unlike deterministic automation, agents reason probabilistically, making judgment calls based on the context available to them. That makes the quality, freshness, and governance of that context the difference between an agent that adds value and one that creates risk.

When organizations build agents as isolated projects, each agent typically handles its own data access, tooling, and security. This approach creates serious problems at scale.

Data connections are duplicated. Security policies are applied inconsistently. And there is no centralized way to audit agent behavior across the enterprise. Disconnected agents also expose an important data latency gap. Trustworthy agentic operations depend on fresh data. A standard two-hour batch refresh between transactional systems and a data platform breaks an agent's closed-loop workflow. When an agent acts on stale inventory or account data, its autonomy becomes a liability.

These are the same structural problems that drove the adoption of service-oriented architecture and modern API-led integration platforms in the first place. Agents introduce a new type of participant in the enterprise architecture, but the principles of good connectivity design still apply.

Gartner estimates that over 40% of agentic AI projects may be canceled by the end of 2027, largely because organizations are building agents without the connective infrastructure to support them in production.



The case for a federated, data-foundation-first connectivity platform

Traditional enterprise software coupled business logic to a specific user interface. Headless architecture has broken that coupling, exposing capabilities as APIs that any front end, any channel, and any consumer could access. Agentic AI completes that evolution. The "front end" is now an agent runtime, and dozens of commercial and custom-built options are already on the market. Organizations should now enable any agent to discover, reason over, and act on enterprise capabilities. APIs and integration become the connective tissue that makes this possible. The enterprise that invested in headless architecture to serve mobile, web, and collaborator channels is already partway there. Agents are the next channel.

Agents perform better and cost less to operate when they are built on top of a well-designed integration layer rather than assembled around ad hoc connections. A federated, domain-owned connectivity model can give agents governed access to enterprise data through cataloged, discoverable interfaces, rather than requiring each agent to build and maintain its own connections.

Organizations should build toward an agentic connectivity platform that sits alongside existing iPaaS, with the Model Context Protocol (MCP), a standardized protocol that lets AI agents discover and connect to any enterprise tool or data source, as the integration standard for AI agents.




This federated approach mirrors a principle central to Deloitte's broader convergence architecture point of view: central control alone slows innovation, and pure domain autonomy creates chaos. The answer in both cases is the same: domains move fast while the enterprise provides shared standards, identity, governance, and reusable connectivity.


The platform connects agents to the tools and services they need through MCP, coordinates communication between agents through the Agent2Agent (A2A) protocol (the open standard for how agents from different systems securely communicate and coordinate tasks, originally developed by Google and now governed by the Linux Foundation), and provides observability built for agentic workloads.

Figure 1
Federated agentic connectivity platform

BEFORE: Fragmented, disconnected environment



AFTER: Connected, governed environment





Each of these architectural requirements maps to a live capability. In the MuleSoft ecosystem, this is no longer a roadmap item. MuleSoft Agent Fabric delivers these capabilities today: an Agent Registry that catalogs each agent and MCP server for runtime discovery, an Agent Broker that organizes agents and tools into business-focused domains and routes tasks across them, Agent Governance built on the Omni Gateway (formerly Flex Gateway) and AI Gateway that applies one set of security and compliance policies to each agent-to-agent and agent-to-tool interaction over MCP and A2A, and Trusted Agent Identity, which carries a user's identity through the agent network so each action stays authorized and auditable. For MuleSoft customers with existing APIs already built, the path is direct: MCPify those APIs, register them in the Agent Registry, and make them discoverable by agents in the ecosystem, an accelerant, not a rip-and-replace.

MuleSoft supports agentic connectivity by turning fragmented systems into reusable, governed APIs that make enterprise data and processes easier for agents to access and act on. Combined with Deloitte Digital's solutions such as Cloud API Vault and AI Assist, organizations can speed up API discovery, reuse, and delivery while strengthening governance at scale.

Traditional monitoring captures logs, metrics, and traces to show what happened in a system. Agentic observability goes further, recording what options an agent considered, what decision it reached, and how confident it was in that decision. MuleSoft's Agent Visualizer is built for this, surfacing confidence scores, bottleneck detection, and hallucination risk across the agent network. This kind of audit trail is essential for compliance, for debugging, and for building organizational trust in autonomous systems, and it maps directly to two dimensions of Deloitte Digital's Trustworthy AI™ framework: transparent and explainable, and accountable. The rise of autonomous agents is a reason to strengthen these established trust dimensions rather than write a separate rulebook.

Central control alone will slow innovation. Pure domain autonomy will create chaos. The winning model is federated: domains move fast, while the enterprise provides shared standards, identity, governance, observability, and reusable connectivity. This platform should follow that federated, domain-owned model, applying the principles of Data Mesh to MCP server deployment. The teams closest to the data, in functions such as supply chain, finance, and HR, are best positioned to build and operate their own MCP servers, while a central platform team provides shared infrastructure, tooling, and governance guardrails.

A federated architecture also needs to distinguish between analytic and transactional MCP servers:



Analytic MCP servers (read and query) pull from cloud data warehouses and analytics platforms. These carry low risk, since they only expose data, and they are a good place to start proving value, provided the underlying data is trusted and governed. This is where platforms like Informatica do the foundational work, profiling, cleansing, and governing data before it ever reaches an agent.



Transactional MCP servers (read, write, and execute) connect to core processing platforms such as Salesforce and enterprise ERP and HR systems. These require real-time, ACID-compliant data, strict role-based access control, and human-in-the-loop workflows for critical operations.





The MCP catalog: The connective tissue of the agent ecosystem

To make a federated model governable and discoverable, organizations need an MCP catalog. Instead of hardcoding connections, agents query the catalog at runtime through semantic search, finding approved tools by matching them to the capability they need.

The catalog is the single enforcement point for security. New servers go through admin approval workflows that prevent shadow AI from proliferating. It maps dependencies to underlying APIs, so integration teams understand the blast radius when source systems change, and it captures broad audit trails (which agent, which tool, what result) that can help with regulatory compliance.

Agentic systems need cross-functional ownership

Designing and operating a full agentic system typically reaches beyond any single team. Integration teams manage the connective layer, including communication protocols, permissions, security, and access control. AI teams build the reasoning models and logic that power agent decisions. UX teams design how humans interact with agents. Data teams maintain the quality and accessibility of the information agents consume. Security and governance teams define what agents are allowed to do, which should be treated as design inputs for the system rather than constraints applied after the fact.


Agentic connectivity is only as strong as the data foundation underneath it. Every capability described so far, including the MCP catalog, federated servers, and agent orchestration, depends on the quality and governance of the data those agents ultimately reach. This is where a shared data foundation becomes strategic. Trusted, enterprise-scale agents draw on four complementary capabilities working together:

- Salesforce Data 360** provides the business context and semantic layer, unifying customer and operational data so agents understand what it means.
- MuleSoft** provides governed connectivity: API reuse, action, and agent-to-agent orchestration.
- Informatica** provides trusted, governed, high-quality data underneath.
- Deloitte Digital** brings the industry insight, customer domain expertise, and operating model experience needed to turn these capabilities into production systems, grounded in how each sector actually makes decisions, serves customers, and manages risk.

Organizations that assign agentic AI to their AI team alone are likely to build solutions that work in a controlled demo but struggle in production, where agents run into real enterprise systems, inconsistent data quality, and regulatory requirements. Getting this right takes cross-functional ownership from the start.

Figure 2

The Data Foundation: Four capabilities



Deloitte Digital's reference architecture for agentic connectivity

AGENT ORCHESTRATION

Agent orchestration coordinates multiple agents on complex tasks, managing dependencies and routing work through standardized protocols such as A2A, so agents from different vendors and internal teams can interact securely and predictably.

THE MCP GATEWAY AND SEMANTIC LAYER

These serve as the centralized access point, combining the MCP catalog, an identity gateway (OAuth 2.1 with PKCE), and dynamic semantic search for runtime tool discovery.

FEDERATED MCP SERVERS

Federated MCP servers are domain-owned connections (REST, gRPC, or CDC) to underlying business data, operated by the teams closest to that data. The architecture distinguishes two types: analytic servers (read-only queries to cloud data warehouses and analytics platforms, low-risk and ideal for an initial deployment) and transactional servers (read-write access to core enterprise systems such as Salesforce and ERP platforms, requiring strict RBAC and human-in-the-loop approval for important operations).



DUAL-PATH DATA ACCESS AND EVENT STREAMING

These route agent reads intelligently, going to source systems for recently modified data and to analytics platforms for historical queries. Change data capture (CDC) on important transactional systems reduces replication lag from hours to seconds.

OBSERVABILITY AND AGENT OPS

These track agent performance, decision quality, and compliance in real time, supporting both operational monitoring and regulatory auditability.

Trusted data and integration foundations, powered by platforms like Informatica, provide the connective tissue across enterprise systems. These capabilities help ensure agents can access unified, governed, and high-quality data while enabling secure integration across cloud, on-premise, and third-party environments.

Getting this right depends less on any single capability and more on the platform and the partner behind it. A trusted platform provides the rails. A partner who has done this before knows how to run on them.



Getting started

For CIOs and CTOs evaluating their readiness, the first step is an honest assessment of the current integration foundation. Rolling out agentic integration works best as a use-case-driven, domain-progressive journey rather than a big-bang deployment. To start:

01

ESTABLISH THE MCP PLATFORM TEAM

Stand up a small platform engineering team to build the shared MCP infrastructure: catalog, gateway, and identity integration.

02

SELECT A PILOT DOMAIN

Identify a single business domain with strong executive sponsorship and high API maturity. Build three to five pilot MCP servers, starting with analytic, read-only use cases, in a two-week discovery sprint.

03

DRAFT ENTERPRISE STANDARDS AND FORM A COUNCIL.

Author minimum viable standards for authentication, logging, and schemas. Convene domain owners, platform leads, and security teams into a governance council.

04

ADDRESS THE LATENCY GAP EARLY

Put dual-path MCP routing in place as a Phase 1 deliverable, and begin evaluating CDC for a mid-term event-streaming solution.

Organizations with fragmented integration layers, inconsistent API governance, or heavy reliance on point-to-point connections may find that agents amplify those problems before they deliver value. Organizations that have already invested in API-led connectivity, strong data governance, and centralized observability are better positioned; for them, agentic capabilities represent a natural extension of work already underway. For organizations earlier in their integration maturity, the investment case is even stronger. Building a sound integration foundation now solves existing connectivity problems while preparing the enterprise for agentic workloads. A practical way to start is to stand up read-only analytic MCP servers against cloud data warehouses and analytics platforms to prove value with low risk, then extend to transactional systems such as Salesforce and core ERP platforms where state changes demand full RBAC and human-in-the-loop approval.

Before scaling agents, leaders should ask one question: are we building autonomous intelligence on top of trusted enterprise foundations, or are we creating a faster, less visible version of the fragmentation we spent the last decade trying to fix?

Deloitte Digital's Data Foundation practice covers the full agentic-connectivity lifecycle, from readiness assessments and federated MCP and reference-architecture design through build-out, security, and ongoing AgentOps. Through its alliances across Salesforce Data 360, MuleSoft, and Informatica, and as co-producer of MuleSoft's annual Connectivity Benchmark Report, Deloitte Digital pairs deep data and integration engineering with the governance discipline of its Trustworthy AI™ framework, helping organizations move agentic AI from proof of concept to production at enterprise scale.

Get in touch

AUTHORS

Kurt Anderson

Salesforce Data Foundation Leader
Deloitte Consulting LLP
kuanderson@deloitte.com

Tushar Sahney

US MuleSoft Practice Leader
Deloitte Consulting LLP
tsahney@deloitte.com

Karthik Santhanagopalan

API and Integration Leader
Deloitte Consulting LLP
ksanthanagopalan@deloitte.com

Prathamesh Kubal

US Informatica Practice Leader
Deloitte Consulting LLP
pkubal@deloitte.com

Karim Trojette

API and Integration Leader
Deloitte Consulting GmbH
ktrojette@deloitte.de

Nick Ingham

US Data 360 Practice Leader
Deloitte Consulting LLP
niingham@deloitte.com

ADDITIONAL READING

Enterprise AI Convergence Architecture, Deloitte AI Institute

<https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/articles/enterprise-ai-convergence-architecture.html>

The State of AI in the Enterprise, Deloitte AI Institute

<https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html>

The 2026 Connectivity Benchmark Report, Deloitte Digital & MuleSoft

<https://www.mulesoft.com/lp/reports/connectivity-benchmark>

Follow the AI Tokens: How CTOs Can Manage Tokenomics, Deloitte

<https://www.deloitte.com/us/en/services/consulting/articles/future-of-enterprise-it-tokenomics-insights-for-cto.html>

Tokenomics: A Guide to Governing the AI P&L, Deloitte via WSJ

<https://deloitte.wsj.com/riskandcompliance/tokenomics-a-cfos-guide-to-governing-the-ai-p-l-ea09aed4>

Trustworthy AI, Deloitte AI Institute

<https://www.deloitte.com/us/en/issues/generative-ai/trustworthy-ai.html>

ENDNOTES

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited (DTTL), its global network of member firms, and their related entities (collectively, the “Deloitte organization”). DTTL (also referred to as “Deloitte Global”) and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/about to learn more.

This communication contains general information only, and none of Deloitte Touche Tohmatsu Limited (DTTL), its global network of member firms or their related entities (collectively, the “Deloitte organization”) is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser. No representations, warranties or undertakings (express or implied) are given as to the accuracy or completeness of the information in this communication, and none of DTTL, its member firms, related entities, employees or agents shall be liable or responsible for any loss or damage whatsoever arising directly or indirectly in connection with any person relying on this communication.