




The Foundation of the Agentic Enterprise

Why integration platforms and patterns
are critical for scalable, governed agentic AI

Executive summary

Enterprises are moving quickly from AI experiments to AI-enabled operations, from customer service and sales to procurement, finance and software delivery. The challenge is whether organisations can scale AI adoption safely while delivering measurable value.

When agentic initiatives stall, the cause isn't typically model capability, but the absence of an operating foundation that makes AI reliable, governable and economically predictable across the enterprise.

The **agentic enterprise** – where autonomous agents execute end-to-end processes across front-office and back-office platforms – promises step-change improvements in cycle time, cost-to-serve and decision quality. But it also introduces new enterprise risks: probabilistic behaviour, opaque decision paths and machine-speed error propagation. So, how do we turn agentic AI into a controlled business capability rather than a collection of pilots?

A second challenge is emerging just as quickly: [the economics of AI](#). Agentic workloads consume tokens continuously through prompts, tool calls, retries, memory, retrieval and monitoring. These costs can be volatile, driven by usage spikes, model selection, context length and error loops. As a result, AI is shifting from a discussion about technology to a CFO and board-level issue. Cost predictability, unit economics (cost per case/order/claim), guardrails on spend, and governance that ties consumption to outcomes are all factors that must be considered. Without deliberate workload design and cost controls, agents can become financially ungoverned.

The scalable path is to **treat agentic AI as enterprise transformation enabled by platforms and patterns**. Specifically, organisations need a robust foundation for discoverability, interoperability, orchestration, observability, security and governance across an increasingly diverse landscape of software as a service (SaaS), data platforms and legacy systems. Integration platforms and patterns extended for an agentic world are a practical place to establish that foundation because they already straddle business processes and systems of record.

Critically, a foundation layer is no substitute for modernisation. Senior leaders must make explicit choices about where to modernise core systems (because process, data, and controls must change at the source) versus where to wrap and extend legacy using application programming interfaces (APIs), integration and emerging agent/tool protocols (because the system is stable, high value and not yet economical to replace).

Over-rotating to 'wrap everything' creates hidden constraints such as data quality limits, brittle semantics and operational risk, capping AI outcomes. The most effective programs combine selective core modernisation with well-governed integration.

Our intent is not to promote a single vendor approach, but to describe what capabilities organisations need to operationalise agentic AI. Commercial platforms (for example, integration and API management suites, eventing platforms, workflow/orchestration engines and agent/tool registries) can accelerate delivery when they implement these capabilities coherently. Vendor offerings like Mulesoft from Salesforce can be viewed as reference implementations of this broader capability set.

1 *The rise of the agentic enterprise*

The agentic enterprise is not a destination, but a framework built on solid architectural foundations. As AI agents proliferate across the enterprise, the absence of stringent governance transforms their evolutionary potential into a catalyst for exponentially compounding chaos.

1.1 THE PARADIGM SHIFT

Agentic AI represents a fundamental shift from isolated AI capabilities to autonomous, goal-driven agents that:

- **Reason and plan** across multiple steps and decision points
- **Interact dynamically** with humans, systems, data sources and other agents
- **Execute tasks** across the front office (including customer relationship management and customer service) and back office (ERP, finance, human resources)
- **Adapt continuously** based on feedback, context, environment and changing conditions.

1.2 THE ENTERPRISE VISION

Modern enterprises increasingly demand the ability to:

- **Plug and play** leading SaaS platforms alongside legacy systems
- **Introduce agents** without extensive custom integration, whether vendor-provided or bespoke
- **Seamlessly unify** human, digital and AI workers through new and refined processes
- **Swap or evolve agents** without destabilising core systems or requiring complete re-architecture.

1.3 THE REALITY WITHOUT FOUNDATION

Without a robust architectural foundation, this vision quickly collapses. Organisations grapple with brittle point-to-point integrations that break with every change, opaque decision-making where no one understands why an agent took a specific action, and ungoverned automation with compliance, security and operational risks. Automated chaos, not intelligent acceleration.



2 *Bridging the chasm: Transitioning AI from pilot success to platform scale*

2.1 THE PROOF-OF-CONCEPT GRAVEYARD

Most organisations can successfully deploy agentic AI in controlled, isolated environments. But their proof-of-concept (PoC) graveyard quickly fills with projects that showed promise but lacked a scalable foundation.

Some projects fail when attempting to span multiple systems and domains, such as connecting sales agents to inventory, finance and logistics. Others may fall over when supporting critical business processes where errors have real consequences, such as maintaining security, compliance and auditability in regulated industries. Similarly, measuring return on investment (ROI) by observing and optimising agent behaviour over time is difficult to scale and often results in failure.

2.2 COMMON FAILURE PATTERNS

Organisations encounter predictable symptoms when scaling without proper foundations. These include:

- **Hard-coded integrations** between agents and applications that cannot adapt
- **Agent-to-agent integrations** need new patterns and considerations
- **Lack of visibility** into agent decisions, reasoning paths and execution flows
- **Inconsistent identity, security and data handling** across different agent implementations
- **No clear ownership or governance model** for agent lifecycle management
- **Infinite feedback loops** where agents trigger each other in costly, uncontrolled cycles.

2.3 THE HISTORICAL PARALLEL

These challenges aren't new. Enterprises have faced similar predicaments in the past, including with service-oriented architecture (SOA) in the 2000s, microservices in the 2010s, SaaS integration and API management, and event-driven architectures and real-time data streaming.

The difference now is that agents don't just respond, they act autonomously. This amplifies both the value and the risk, intensifying the need for robust foundations.

3 *Integration platforms as the foundation for agentic AI*

Integration platforms are uniquely positioned to underpin the agentic enterprise because they already address the hardest problems at enterprise scale. The evolution from traditional integration to agentic foundations requires extending five core capabilities.

3.1 INTEROPERABILITY: THE UNIVERSAL CONNECTOR

The challenge:

In a plug-and-play world, the biggest hurdle isn't connectivity, it's meaning. For example, if a bespoke research agent asks for 'customer lifetime value', the back-office system might define it as `net_profit_total`, while the front-office CRM defines it as `projected_contract_value`.

This is where modern integration platforms can offer a solution.

- **A semantic translation layer** transforms data into the 'language of business' allowing agents to see what the business means, rather than technical field names
- **Application-to-application (e.g. Connectors)** communication for traditional systems
- **Agent-to-agent (A2A)** interactions for autonomous collaboration with shared context
- **API-based and event-driven patterns** for real-time responsiveness
- **Model context protocol (MCP)** support for emerging standards for tools/skill building.

Everything old is new again. The patterns are familiar – standardised interfaces, reusable building blocks, governance, and observability – but the stakes are higher. Agentic AI moves faster, operates with uncertainty, and incurs usage-based costs at scale.

The organisations that win will be those that build operating foundations to keep agents safe, auditable, and economically governed, so AI adoption translates into durable business value rather than isolated technical success.

KEY ARCHITECTURAL PATTERN:

The agentic control plane vs. integration control plane split

Drawing from networking and cloud architecture, the agentic enterprise should be separated into two separate branches.

1 *Agentic control plane:*

Where governance and policy live. This branch defines which agents are 'certified' for specific tasks and manages model context protocol (MCP) to ensure shared understanding.

2 *Integration plane:*

The 'plumbing' (MuleSoft Fabric, iPaaS) handling protocol translation between modern large language models (LLMs), such as JSON/Natural Language, and legacy/traditional systems.

Interoperability ensures agents are accurate, consistent, composable, replaceable, and evolvable.

3.2 TRANSFORMATION: CONTEXT IS THE CURRENCY OF AGENTS



The challenge:

Raw data from back-office ERPs is rarely 'LLM-ready'. Agents need rich, accurate and contextual information to make decisions, and enterprise data is often fragmented, inconsistent, and ambiguous.



The solution:

The semantic translation layer, with integration platforms providing:

- **Data transformation and enrichment**, turning raw formats into content that agents understand
- **Normalisation across disparate systems**, ensuring consistent business semantics
- **Schema and contract management** for reliable, versioned data exchanges
- **Semantic gateway functionality** that transforms data into a 'Language of Business'.

This ensures agents operate on:

- **Consistent business semantics** (not technical field names)
- **Trusted and governed data** (with lineage and quality controls)
- **Context appropriate to their role and scope** (sales agents see different data to finance agents)
- **Shared necessary restricted context across AI agents working together** (sales agent needs to update contact preferences in service centre using service agent).

Without this layer, agents inherit the chaos of raw enterprise data, leading to 'hallucination of logic' where agents misinterpret ambiguous fields and make incorrect decisions.

Transformation prevents agent errors and enables accurate, context-aware reasoning.

3.3 ORCHESTRATION: FROM TASKS TO END-TO-END PROCESSES



The challenge:

True business value emerges when agents participate in end-to-end workflows, not isolated tasks. Traditional integration relies on orchestration, where a central brain follows a script. The agentic enterprise moves toward choreography; that is, independent agents reacting to events.



The solution:

Dynamic choreography with guardrails. Integration platforms enable process orchestration across systems and agents for complex workflows, coordination between human and AI workers, long-running, stateful business processes that survive failures and restarts, and event bus architecture for event-driven agent collaboration.

Example:

When a 'shipment delayed' event occurs:

1

The foundation
broadcasts
the event

2

The Logistics
Agent updates
tracking
information

3

The Customer
Service Agent
drafts a proactive
apology email

4

The Inventory
Agent adjusts
reorder schedules

The risk:

Agents trigger each other in infinite, costly cycles with no ability to debug or control the system.

Human-in-the-loop (HITL) integration:

Platforms like MuleSoft Agent Fabric coordinate interactions, ensuring processes don't 'break' when moving from autonomous agents to humans for high-stakes approvals. Standard integration patterns pause agent flows and 'ping' humans via Slack or Teams when risk thresholds are met.

As a result, agents are participants in orchestrated flows rather than uncontrolled actors, ensuring reliability, repeatability, and resilience.

3.4 OBSERVABILITY: SEEING, MEASURING, AND TUNING AGENT BEHAVIOUR



The challenge:

Agentic systems must be observable to be trusted. Enterprise leaders hesitate to move beyond proofs of concept (PoCs) because agents are inherently probabilistic. If agents are unobservable, they can't be regulated, creating hazards in mission-critical environments.



The solution:

Real-time guardrails and auditability. Using a foundational platform gives you end-to-end tracing across systems and agents with complete execution paths, as well as metrics on execution, latency, and outcomes for performance optimisation. Regulatory requirements can be logged for audit, debugging, and compliance, and reasoning traces capture not just API calls, but the 'thought trace' of agents.

The observability tax: To bridge the trust gap, foundations must implement:



Semantic firewall: Inspecting agent outputs before they reach systems of record (e.g., 'An agent cannot offer a discount greater than 20%')



'Why' traces: Logging the reasoning (prompt, retrieved context, tool call) instead of just outcomes



Deterministic fallbacks: If confidence scores drop below thresholds, automatically reroute to standard APIs or human-in-the-loop



Cost & token management: Tracking ROI of each agentic transaction (token cost vs. business value)

Using the observability tax enables enterprises to track agent decisions and actions with full transparency, identify bottlenecks or failure points for continuous improvement, continuously tune and improve performance through feedback loops, and treat agents as a 'digital workforce' with performance reviews.

Ultimately, moving beyond PoCs requires a feedback loop that tunes the 'probabilistic' nature of AI into 'deterministic' business outcomes.

3.5 SECURITY AND GOVERNANCE: NON-NEGOTIABLE IN THE ENTERPRISE



The challenge:

If agentic AI is not governed properly, it amplifies risk. In an agentic enterprise, 'identities' are no longer just for humans. Every agent requires a machine identity with granular permissions.



The solution:

Centralised governance at the integration layer. Integration platforms enforce:

- **Centralised authentication and authorisation** for all agents, actions, tools & skills
- **Policy-based access control** defining what agents can view, modify, or execute
- **Data privacy and residency controls** ensuring compliance with the GDPR, CCPA, Privacy Act etc.
- **Versioning and lifecycle management** for agent evolution and rollback
- **Machine identity implementation** where every agent has its own credentials and scopes.

This allows enterprises to apply consistent governance across all agents regardless of origin and ensure compliance with regulatory requirements through centralised policy enforcement. It also safely enables autonomy within defined boundaries, preventing 'agent sprawl' and unauthorised access to sensitive personally identifiable information (PII) through granular permission models. Autonomy without governance is exposure, not innovation.

Outcome:

Security and governance transform agents from potential liabilities into trusted, governed members of the digital workforce.



4 *Moving beyond pilots: The role of a unified integration fabric*

4.1 THE FOUNDATION-FIRST APPROACH

A unified integration fabric provides a single control plane for integration, automation and AI agents across all systems. It includes consistent patterns across front-office and back-office platforms, reusable building blocks for agents and applications, and a centralised agent registry for discovery and governance.

Reference architecture: Mulesoft Fabric

MuleSoft Fabric serves as an example of the 'foundation-first' approach:



Discover:

Register agents in a central catalogue (Anypoint Exchange) with semantic documentation.



Connect:

Use MCP to bridge agents to any data source, legacy system, or modern API.



Govern:

Apply consistent security policies across all agentic interactions through API Manager.



Observe:

Track agent behaviour, performance, and outcomes through unified observability.



Orchestrate:

Coordinate multi-agent workflows with human-in-the-loop capabilities.

4.2 THE TRANSFORMATION ENABLED

By placing this foundation first, enterprises can scale successful use cases across the organisation and rapidly introduce new agents and platforms without custom integration work. They can also reuse patterns to avoid rebuilding the same integration logic repeatedly, and support end-to-end business processes that span multiple systems and domains.

This step is critical when moving from experimentation to enterprise capability.

4.3 COMPARISON: TRADITIONAL VS. AGENTIC INTEGRATION

FEATURE	TRADITIONAL INTEGRATION	AGENTIC FOUNDATION
Connectivity	API-to-API	Agent-to-agent (A2A) & model context protocol (MCP)
Logic	Static workflows	Dynamic planning/reasoning loops
Data type	Structured	Structured + semi-structured + unstructured context
Security	OAuth/API keys	On-behalf-of (OBO) authentication, machine identity & trust guardrails
Discovery	Developer reads docs	Agent queries MCP catalogue
Orchestration	Central brain (script)	Event-driven choreography
Governance	IP Whitelisting/OAuth	Semantic guardrails/policy-based/OAuth2 with scopes

5 *The roadmap: From API-led to agentic-led architecture*

To transition from API-led to agentic-led architecture, enterprises must shift focus from providing data to enabling action. The following four-phase roadmap outlines the technical evolution required.

Phase 1:

Foundation (the 'agent-ready' API)

GOAL:

Transform passive APIs into machine-interpretable tools.

KEY ACTIVITIES:

Semantic documentation: Move beyond basic Swagger/OpenAPI. Enhance specs with metadata describing intent and pre-conditions (e.g., 'Use this API only if customer has Gold status').

- **Rationalise API layers:**
 - **System APIs:** Ensure access to the original, unfiltered, raw data
 - **Process APIs:** Refactor to be domain-centric (e.g., reconcile_invoice instead of update_record_v2).
- **Security baseline:** Implement machine identity so that every agent has its own credentials and scopes.

OUTCOME:

APIs become discoverable, interpretable tools for agents rather than passive endpoints.

Phase 2:

Interoperability (The Protocol Shift)

GOAL:

Implement model context protocol (MCP) to enable plug-and-play capability.

KEY ACTIVITIES:

- **Deploy MCP servers:** Wrap existing API Gateways or MuleSoft flows in MCP Server layers
- **Standardise data transformation:** Convert complex ERP formats into LLM-friendly JSON or Markdown formats
- **Establish agentic registry:** Create a central catalogue (evolution of API Portal) where agents ask about available capabilities at runtime.

OUTCOME:

Any LLM-based agent (OpenAI, Anthropic, bespoke) can automatically 'see' and use enterprise APIs as tools.

Phase 3:

Control & orchestration (the agentic fabric)

GOAL:

Move from 'if-then' scripts to 'dynamic choreography'.

KEY ACTIVITIES:

- **Implement agentic guardrails:** Place semantic firewalls between agents and back-office systems to inspect intent before execution
- **Build HITL connectors:** Create standard patterns that pause autonomous flows for human approval at high-risk thresholds
- **State management:** Use distributed state stores (Redis, Vector DBs) so supervisor agents can resume failed processes.

OUTCOME:

Controlled, governable agent collaboration with safety mechanisms and human oversight.

Phase 4:

Observability & continuous tuning

GOAL:

Treat agents as a 'digital workforce' with performance reviews.

KEY ACTIVITIES:

- **Capture reasoning traces:** Log not just API calls, but the 'thought trace' (why did the agent choose API A over API B?)
- **Implement cost & token management:** Track ROI of each agentic transaction through integration layer headers
- **Build feedback loops:** Feed success/failure data back into agent prompt context for self-correction.

OUTCOME:

Continuous improvement, measurable ROI, and transformation of probabilistic AI into deterministic business outcomes.

6 *Everything old is new again*

6.1 THE PARADOX OF PROGRESS

The agentic enterprise is designed to reinforce existing architectural lessons. The core problem of the enterprise has always been disorder. Without robust foundations, agentic AI will not be an accelerator; it will be a source of automated chaos.

6.3 THE AMPLIFICATION EFFECT

The autonomy and speed of agents magnify both value and risk. Value is amplified as agents can execute in seconds what took humans hours, but risk is equally amplified as ungoverned agents can cause damage at machine speed.

These might be familiar patterns, but the stakes are higher now, and the need for foundations is absolute.

6.2 LESSONS FROM INTEGRATION HISTORY

We know that:

- **Point-to-point integrations do not scale**
(SOA taught us this)
- **Governance must be designed, not bolted on**
(microservices reinforced this)
- **Observability is essential for trust**
(DevOps made this clear)
- **Platforms outperform projects**
(cloud transformation proved this).

6.4 THE API-FIRST MOVEMENT'S FINAL FORM

The agentic enterprise is effectively the final form of the API-first movement. If your enterprise does not have a mature API strategy, your agents will be 'blind' with no hands to reach into systems and no eyes to see data.

By building on an integration-centric foundation, you ensure agents are not siloed bots but integrated members of the digital workforce that are interchangeable, governable, and scalable.



7 *Foundations first, agents second*

7.1 THE DEFINING FACTOR

The success of the agentic enterprise will not be determined by who has the most sophisticated models or the flashiest agents. It will be determined by who has built the strongest foundations.

These organisations will be able to:

- Move confidently from PoCs to production
- Scale agentic capabilities across the enterprise
- Maintain security, compliance, and governance
- Deliver measurable, sustained business outcomes.

7.3 THE UNIVERSAL FABRIC VISION

Whether you are deploying a Salesforce agent for customer service, a custom-built Python agent for data analysis, a legacy SAP bot for procurement or a vendor-provided agent for HR onboarding, they all must adhere to the same enterprise standards of security, reliability, and governance. This is only possible with a universal integration fabric as the foundation.

The future belongs to the organisations that recognise agents are only as powerful as the platforms they stand upon.


7.2 THE PATH FORWARD

By extending integration platforms to support agentic patterns – including discoverability, interoperability, transformation, orchestration, observability, security, and governance – enterprises can confidently move beyond experimentation toward scalable, governed, and business-critical agentic systems.



Get in touch

Contact us to learn how a strong integration foundation can help you scale agentic AI safely and confidently.



Paul Fayle

Partner, Technology
& Transformation
Deloitte Consulting LLP
pfayle@deloitte.com.au



Sasi Sasindran

Engineering Director,
Technology & Transformation
Deloitte Consulting LLP
ssasindran@deloitte.com.au

Deloitte. Digital

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 organisation') 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.

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 organisation'). 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.

About Deloitte

Deloitte is a leading global provider of audit and assurance, consulting, financial advisory, risk advisory, tax and related services. Our global network of member firms and related entities in more than 150 countries and territories (collectively, the 'Deloitte organisation') serves four out of five Fortune Global 500® companies. Learn how Deloitte's approximately 400,000 people make an impact that matters at www.deloitte.com.

About Deloitte Asia Pacific

Deloitte Asia Pacific Limited is a company limited by guarantee and a member firm of DTTL. Members of Deloitte Asia Pacific Limited and their related entities, each of which are separate and independent legal entities, provide services from more than 100 cities across the region, including Auckland, Bangkok, Beijing, Bengaluru, Hanoi, Hong Kong, Jakarta, Kuala Lumpur, Manila, Melbourne, Mumbai, New Delhi, Osaka, Seoul, Shanghai, Singapore, Sydney, Taipei and Tokyo.

About Deloitte Australia

The Australian partnership of Deloitte Touche Tohmatsu is a member of Deloitte Asia Pacific Limited and the Deloitte organisation. As one of Australia's leading professional services firms, Deloitte Touche Tohmatsu and its affiliates provide audit, tax, consulting, risk advisory, and financial advisory services through approximately 12,000 people across the country. Focused on the creation of value and growth, and known as an employer of choice for innovative human resources programs, we are dedicated to helping our clients and our people excel. For more information, please visit our web site at <https://www2.deloitte.com/au/en.html>.

Liability limited by a scheme approved under Professional Standards Legislation.

Member of Deloitte Asia Pacific Limited and the Deloitte organisation.

© 2025 Deloitte Touche Tohmatsu

14742795_Designed and produced by Brand and Marketing | Deloitte Australia_06/26