# The Agentic AI Stack

## Reconsidering vendor choices in the era of autonomous AI

A whitepaper by newITera
2026

| 62% | 29% | 78% | 9% |
|---|---|---|---|
| cite complexity as the main barrier (KPMG, 2026) | already use unsanctioned AI agents (Microsoft, 2026) | surprised by unexpected AI pricing models (Zylo, 2025) | are fully AI-ready on data (Huble, 2025) |

# Executive Summary

## The stack as a whole, or not at all

Agentic AI, systems that reason independently, make plans, and execute tasks, stands on the threshold of broad enterprise adoption. But anyone who thinks they can get by with a separate tool choice per department fundamentally underestimates what is at stake.

The core message of this whitepaper is simple: an application choice is no longer just an application choice. It is a choice for a complete technology architecture of five layers, from infrastructure to the interface your employees work with every day. Every layer is moving fast and influencing the others. Those who fail to recognize this will soon find themselves on tracks a vendor has laid for them.

This whitepaper consolidates insights from proprietary market research, platform analyses of nine leading vendors, and recent studies from KPMG, Info-Tech Research Group, Microsoft, and the FinOps Foundation. It offers enterprise organizations an objective evaluation framework to weigh AI platform choices at a strategic level, with particular attention to organizations at the beginning of their AI journey.

**Four messages you will take away from this whitepaper:**

- Vendor lock-in is no longer a theoretical risk, but an architectural reality being determined right now.
- The AI stack consists of five layers you need to evaluate as a portfolio, not as separate products.
- Your AI maturity determines which entry point fits your organization.
- newITera helps you navigate a vendor landscape that has never been busier or more strategically consequential.

# The problem: three new pressures on your technology decisions

Organizations are experimenting with AI on a massive scale. Standalone tools are being rolled out, pilots are being launched, and dashboards are appearing everywhere to demonstrate AI use. But structural impact remains elusive. The gap between AI hype and operational value is real, and the explanations are equally real.

Info-Tech Research Group identifies three fundamental pressures that characterize the era of Agentic AI, and that together put the technology decisions of every enterprise organization under tension.

## Technical complexity is stalling implementation

62% of organizations cite the complexity of agentic systems as the primary barrier to deployment (KPMG, 2026). Agentic AI is not a chatbot you configure. It is systems that reason, call tools, maintain memory, and take autonomous steps. That requires an architecture that holds together from top to bottom.

### Governance is lagging behind usage

While IT departments are still writing policy, 29% of employees have already used unsanctioned AI agents for work tasks (Microsoft, 2026). Even more alarming: only 8% of organizations state that their agents never exceed intended permissions (Cloud Security Alliance, 2026). Shadow AI is no longer a marginal phenomenon. It is the reality in your organization, right now.

### Pricing models have become unpredictable

42% of SaaS vendors have switched to variable AI pricing based on usage or outcome (High Alpha, 2025). 78% of IT leaders report unexpected costs from consumption-based AI pricing models (Zylo, 2025). With Agentic AI this problem is larger than with regular AI: a single user instruction can trigger dozens of billable events, from tool calls and model escalations to recursive reasoning loops and repeated API calls.

These three pressures are compounded by a vendor landscape that has never been busier. Leading vendors such as Microsoft, Salesforce, ServiceNow, and SAP are aggressively competing to occupy the central position in your AI architecture. The stakes: whoever owns your agentic control plane determines your IT freedom for the next three to five years.

> *The price of the right vendor choice you pay today in attention. The price of the wrong vendor choice you pay over the coming years in dependency.*

## The five layers of the Agentic AI stack

### Every layer is a moving target

The enterprise AI architecture consists of five inseparably connected layers. The central insight is this: with Agentic AI, you need to oversee all five layers simultaneously. An application choice has implications for your data architecture. A model choice has implications for your infrastructure costs. A platform choice has implications for your governance freedom.

| # | Layer | Role | Strategic development |
|---|---|---|---|
| 1 | **Applications** | *The agentic control plane* | Applications are becoming orchestration platforms. Whoever owns this layer determines your IT freedom for the next 3–5 years. |
| 2 | **Data & AI Tools** | *Orchestration & integration* | Interoperability standards (MCP, A2A) are maturing. Vendors embracing open standards vs. those building proprietary ecosystems. |
| 3 | **Foundation Models** | *The brains* | Multi-model is the enterprise norm. 65% combine multiple models. Dynamic routing dramatically reduces costs. |
| 4 | **Data Platforms** | *Context & meaning* | Data storage is shifting to context delivery. Only 9% of organizations are fully AI-ready on data. |
| 5 | **Infrastructure** | *Compute & governance* | Agentic AI makes traditional budgeting impossible. AI FinOps is no longer optional, it is baseline infrastructure. |

### Layer 1: Applications, the battle for your agentic control plane

The top layer is the strategic battleground. Software giants are rebuilding their applications from transactional systems into agentic orchestration platforms. Microsoft has Copilot and Agent 365. Salesforce launched Agentforce with a headless API architecture and Agent Fabric. ServiceNow

introduced Otto and its AI Control Tower. SAP Joule is evolving into a multi-agent orchestrator, and with SAP AI Agent Hub, agents can be managed centrally.

Behind this terminology lies a concrete strategic reality: whoever owns the application layer owns your architecture. The key question you must ask: who owns the orchestration and governance in my application layer, and what does that dependency look like in three to five years?

## Layer 2: Data & AI Tools, interoperability as the minimum

One layer down, vendors determine how agents communicate with each other. The market is in full motion. The Model Context Protocol (MCP), an open standard that delivers context to language models in a standardized way, is gaining ground. Agent-to-Agent (A2A) communication, the Agent Network Protocol (ANP), and the Agent Client Protocol (ACP) are all attempting to standardize communication between heterogeneous agents.

But beware the subtle distinction: interoperability is the floor, not the ceiling. Virtually every major vendor claims to be open. Salesforce, Microsoft, and Oracle state that their agents communicate with all agents, but add that they perform best within their own ecosystem. Open standards offer porousness; they offer no guarantee of effectiveness outside the vendor perimeter.

## Layer 3: Foundation Models, multi-model as the norm

65% of enterprise organizations now combine two or more commercial models with at least one open-source model (F5, 2026). Multi-model is no longer the exception but the standard. The reason is simple: it is economically illogical to deploy the same model for a routine summary as for complex legal analysis or advanced coding.

Dynamic model routing, where tasks are automatically directed to the most suitable and cost-efficient model, is becoming a baseline function of any serious AI architecture. Alongside proprietary models from OpenAI, Anthropic, and Google, open-weights models such as DeepSeek and open Alibaba models are gaining ground for use cases where sovereignty and privacy on the organization’s own perimeter are required.

## Layer 4: Data Platforms, from storage to context provider

Agents are blind without context. Only 9% of organizations are fully AI-ready on data, while 29% describe themselves as only moderately prepared (Huble, 2025). Data platforms are therefore undergoing a fundamental transition: from systems that store data to systems that deliver meaning.

Agents do not need rows and columns; they need semantic context: what does this data point mean in the context of this business process, what authorizations apply, where does this data come from, and how reliable is it? That semantics, lineage, and governance must travel with the data to the agent. Otherwise your agent is making decisions on wrong or incompletely understood information.

## Layer 5: Infrastructure, AI FinOps as a required discipline

Agentic AI has a fundamentally different cost profile than traditional software. Where costs used to scale with the number of users or a fixed license, they now scale with execution depth. A single user instruction can trigger a chain of tool calls, model escalations, recursive reasoning loops, background API calls, multi-agent orchestration, and retries on errors. Every step is a billable event.

AI FinOps, the discipline of financial management of AI expenditure, thereby becomes a required component of your AI infrastructure. In 2026, 98% of global FinOps professionals are responsible for managing AI expenditure, compared to 31% in 2024 (FinOps Foundation, 2026). Real-time monitoring that consolidates cloud compute and external SaaS token costs is not a luxury but baseline infrastructure to prevent costly runaway loops.

# The vendor landscape: three paths to the future

When an enterprise organization that relies heavily on an SAP landscape decides to automate processes with Agentic AI, the choice of SAP Joule may seem obvious. The hard reality is that an application choice is irrevocably a choice for the entire underlying technology stack. The technology, standards, players, and costs in every layer of the AI stack are changing so rapidly that today's reality may look completely different tomorrow. By overseeing all layers simultaneously today, you retain the agility to adjust tomorrow.

The platform choice fundamentally determines which vehicle you choose for this dynamic journey:

- **The Train (Vendor-Native):** You travel fast and comfortably on pre-laid tracks. The route is fixed. If the market changes, because external models suddenly become more powerful and cheaper, you cannot adjust course.
- **The Car (Decoupled Architecture):** You keep the wheel in your hands. You remain flexible to adjust course, switch models or data sources, or reroute if a better destination appears elsewhere.

This makes the traditional procurement model based on functional product shortlists definitively a thing of the past. Major software vendors are actively competing to become your central agentic control plane. If a vendor succeeds, that vendor will dictate your IT freedom across five critical dimensions: model flexibility, manageability of AI FinOps costs, freedom of your business data, governance over your agents, and your negotiating position at contract renewals.

| THE TRAIN Vendor-native architecture | THE CAR Modular open architecture |
|---|---|
| You travel fast and comfortably on pre-laid tracks. The route is fixed. If external models suddenly become more powerful or cheaper, you cannot adjust course. The vendor determines your destination. | You keep the wheel in your hands. You can adjust course, switch models, switch data sources, or reroute when a better path appears. More responsibility, more freedom. |

The question is not whether the train is bad. For many organizations, the train is exactly the right vehicle right now: fast, safe, and predictable. The question is: are you aware that you have chosen a train, and do you know which track it is taking you on?

## The new evaluation model for the agentic era

The traditional procurement model, based on feature comparisons and point-product shortlists, is definitively a thing of the past. The agentic era demands a fundamentally different evaluation approach.

| OLD EVALUATION MODEL | AGENTIC-ERA EVALUATION |
|---|---|
| *Feature comparison* | Mapping architectural dependencies |
| *Point-product shortlist* | Control plane & ecosystem fit |
| *One-time cost estimate* | Consumption and usage forecast |
| *Vendor demo checklist* | Live agent task evaluation |
| *IT fit & integration check* | Governance, security & lock-in |

Procurement and vendor management teams must now make the transition to architectural thinking. Not: what does this product do? But: how does this platform fit into our five-layer architecture, who owns which layer, and what is the lock-in implication over five years?

## What does this require from your organization?

Agentic AI introduces three types of organizational requirements that each deserve separate attention, but which in practice are closely intertwined.

### Technical prerequisites

The foundation is data readiness. Only 9% of organizations are fully AI-ready on data. That means for 91%: before Agentic AI can deliver structural value, foundational work needs to be done. Think about data quality and consistency, semantic enrichment of business data so agents understand what they see, lineage and governance that travel with the data, and secure connectivity between data sources and agent runtime.

In addition, a clear choice is needed for the orchestration architecture: which platform coordinates the agents, via which protocols do they communicate (MCP, A2A), and how is model routing designed for cost optimization?

### People and change: roles that are shifting

Agentic AI is fundamentally shifting roles. Employees who previously executed tasks become supervisors of agents that execute those tasks. This requires AI Fluency at all levels of the organization: the ability to collaborate with AI systems effectively, efficiently, ethically, and safely.

The four core competencies from the AI Fluency framework apply directly to the Agentic AI context: Delegation (what do you give to the agent, what do you keep yourself?), Description (how do you give an agent the right instructions?), Discernment (how do you evaluate what an agent returns?), and Diligence (how do you work responsibly and transparently with autonomous systems?).

### Governance and accountability

Shadow AI is not a future risk. 29% of employees are already using unsanctioned AI agents. This means governance is needed now, not once the pilot is successful. Effective AI governance for Agentic AI includes at minimum: a clear authorization matrix (who may deploy which agent for which task), audit trails that capture the full reasoning chain of agents, human checkpoints for high-risk decisions (human-in-the-loop), and EU AI Act compliance for high-risk applications.

**Governance reality check:**

Ask yourself: if tomorrow an agent in your organization makes an incorrect decision that leads to an error in a customer order, an incorrect financial report, or an unintended data change, who is responsible? Can you reconstruct the agent's decision? Should a human have stopped this? If the answer to any of these questions is unknown, your governance is not yet ready for Agentic AI.

## How to get started: a phased approach, not a big bang

The most common mistake in enterprise AI initiatives is the opposite of what you might expect: not too little ambition, but too much at once. Organizations that try to build the full five-layer architecture in a single program get stuck on the complexity. The organizations that succeed start small, learn quickly, and build deliberately.

The core strategic advice for organizations with low to moderate AI maturity: secure governance centrally, but start execution vertically. Begin with manageable, safe, and quickly-yielding use cases within existing application silos. Build iteratively from there toward cross-silo orchestration.

### Phase 1: AI Readiness Assessment (weeks 1–4)

Before choosing a platform, you need to know where your organization stands. A thorough AI Readiness Assessment maps out: data maturity, current governance structure, AI skills of your people, legal and compliance context (GDPR, EU AI Act), and the use cases with the highest strategic value and lowest implementation complexity. This assessment forms the objective basis for all platform choices that follow. Without this foundation, you are driving blind.

### Phase 2: Use case selection and validation (weeks 5–8)

Based on the assessment, select one to three concrete use cases that meet three criteria: high business value, low implementation risk, and technical feasibility within your current data maturity. Validate the business case, including a realistic ROI calculation with attention to hidden costs such as API limits, data egress fees, and consumption-based licensing models.

### Phase 3: Pilot and guided implementation (weeks 9–20)

The selected use cases are implemented as a pilot within the appropriate architecture layer. Human supervision is mandatory in this phase: agents act, people evaluate and adjust. The pilot delivers not only technical insight but also organizational lessons about change management, training, and governance.

### Phase 4: Scaling and portfolio thinking (from week 21)

Successful pilots are carefully scaled. Simultaneously, the five-layer architecture is deliberately built based on the insights gained. Vendor evaluations are now conducted at the portfolio level, with explicit attention to lock-in risks, model flexibility, and cost predictability.

> *The organizations that demonstrably perform better in three years thanks to Agentic AI are not those that chose the biggest platform. They are those that started deliberately small, drew the right lessons, and then scaled with confidence.*

## Conclusion: four questions you should be able to answer tomorrow

Agentic AI is not a future technology. It is the current reality in your vendor landscape, in your software contracts, and in the behavior of your employees, who are deploying AI agents, whether sanctioned or not. The question is no longer whether your organization will encounter this. The question is whether you are approaching it deliberately.

| | |
|---|---|
| **1. Control plane clarity**<br>Who owns the orchestration and governance in your application layer, and what does that dependency look like in three to five years? | **2. Interoperability posture**<br>Does your vendor embrace open standards like MCP and A2A, or is it building a proprietary architecture that keeps you locked in? |
| **3. Model & data portability**<br>Can you deploy multiple models and move your data context without significant reconfiguration costs? | **4. Economic transparency**<br>Can your vendor show a predictable cost model for agentic workloads, including retries, tool calls, and orchestration depth? |

If you cannot answer one or more of these questions with confidence, that is not a weakness. It is the right starting point for a conversation. newITera helps you answer those questions, with an objective lens, practical experience, and a phased approach that fits where your organization stands today.

**Ready to take the first step?**

Get in touch for a no-obligation conversation about an AI Readiness Assessment for your organization. We map out where you stand, which use cases offer the most return, and how to future-proof your vendor architecture.

Contact us at laurens.vandebergh@newitera.nl or visit www.newitera.nl

## About newITera

newITera is a Dutch IT consultancy based in Beuningen, founded in 2011. What began as a specialized SAP partner has grown into a broad digital transformation partner with expertise in SAP, data analytics, AI, cloud strategy, and low-code platforms. We combine technical depth with practical insight and a people-centered approach.

Our AI services consist of four pillars: AI Readiness Assessments that measure your organization's maturity at the technical, process, and human level; use case realization where we guide the full journey from selection to implementation; validation and governance where we help structure the assessment of AI output and ensure responsible deployment; and coaching and adoption programs through which teams and individuals develop AI Fluency as a lasting skill.

Laurens Vandebergh is Business Development Manager AI and Project Manager at newITera. He guides organizations in defining their AI strategy, conducting an AI Readiness Assessment, and realizing concrete use cases in the Agentic AI era.

Contact: laurens.vandebergh@newitera.nl | www.newitera.nl