

HFS

In partnership with



MARKET IMPACT REPORT

Design your agent OS to win the AI future

A field guide from early Orchestrators

Authors:

David Cushman, Executive Research Leader

Niti Jhunjhunwala, Senior Analyst

Phil Fersht, CEO and Chief Analyst

Sam Duncan, Practice Leader

Contents

	Page
Firms must have imagination and foundations to scale AI	03
Action and leadership are the critical differentiators in agentic AI	04
Executive summary	05
Misjudge the rate of agentic AI adoption, and you'll watch the leaders sail over the horizon	06
As complexity compounds, change how you architect and govern AI	07
Deliver enterprise-grade reliability with a five-layer blueprint for an agent OS	12
The seven-step field guide to your agent OS	15
The Bottom Line	16
Appendix	17

Firms must have imagination and foundations to scale AI

Agentic AI is becoming a disruptive force in the enterprise, but most organizations are not yet prepared for the breadth and depth of this shift. The challenge now is whether businesses have the imagination and foundations to scale agentic AI safely and effectively. Many deployments begin with copilot-level, incremental use cases such as robotic process automation, workflow automation, and straightforward productivity gains. These are valid entry points, yet they fall short of the transformative potential multi-agent systems can deliver. At the same time, the speed at which some teams are moving toward more complex agentic problem-solving raises the stakes for trust, safety, governance, and richer, context-specific interfaces that help humans and agents collaborate naturally.

The pace of adoption is striking. Orchestrators, adoption leaders with systems using five or more agents, already average around 12 agents in their most mature and successful deployments, even though practical, scalable agents have only been viable for about a year. Few enterprise technologies have ramped this quickly. One reason is that agents can wrap legacy systems, using natural language to connect disparate tools without forcing organizations to rebuild rigid APIs or outdated user experiences. We have helped many clients do exactly that, and we have seen multi-agent systems ease data-readiness barriers.

Businesses are starting to recognize that agentification is not a turnkey solution and that safe, reliable, effective agentic systems require careful design and contextual customization. You do not just deploy a model; you design a system, define the tools agents can use, set boundaries of autonomy, and engineer the interactions that make agents reliable, governable, and useful at scale.

Scaling these systems brings real challenges and demands an operating model purpose-built for autonomy, coordination, and trust. That is why the report's conclusion is clear: You do not need another agent; you need an agent operating system (OS). A unified agent OS turns scattered pilots into cross-functional automation at scale, with increased trust and reduced operational risk. This report distills what leading Orchestrators are learning into a repeatable path forward. Few enterprises will master this alone. Winners will build their agent OS deliberately, backed by a strong partner ecosystem, and convert fast adoption into lasting, business-transforming impact.

The future will be agentic, and the winners will be those who design for it deliberately, starting today.



Babak Hodjat
Chief AI Officer, Cognizant

Action and leadership are the critical differentiators in agentic AI

As we head into 2026, the era of theorizing and testing is over. The moment now is to act, not by designing yet another agent, but by deliberately designing an agent OS that can scale autonomy with control.

Early adopters have the battle scars to prove what happens when organizations move fast without the right foundations. Agents behave less like software and more like employees. Enterprises encounter emergent behaviors, cascading failures when one agent breaks others, and growing blind spots as visibility erodes with each additional agent. Long-cherished processes, once automated, often expose fragility rather than efficiency. This is not a traditional software rollout, and treating it like one is where many stumble.

The data makes the urgency unmistakable. Orchestrators already run about a dozen agents in production, some with as many as 20, and 24% move from pilot to live in under three months. Technology is moving quickly. Leadership systems are not. The widening gap between what AI can do and how decisively organizations are prepared to govern it is the AI Velocity Gap.

When leadership hesitates, agentic systems expose the cracks. Twenty-two percent (22%) of enterprises experienced emergent behaviors, 21% saw cascading failures, and 18% faced auditability gaps after going live. These are not technical surprises; they are leadership ones. Closing the [AI Velocity Gap](#) demands executives actively design autonomy, accountability, and human-machine collaboration. The enterprises that will lead are not those with the most agents, but those willing to confront leadership debt and build an agent OS that scales intelligence with trust.

The future will be agentic. Leadership will decide who wins.



Phil Fersht
Founder & CEO, HFS Research

Executive summary

Multi-agent AI systems are already in production across Global 2000 enterprises at a pace few predicted. Orchestrators typically run around 12 agents in their most mature deployments, with some systems reaching 20, and 24% of them move from proof of concept (POC) to live deployment within three months. Many have stopped running pilots entirely. They're convinced the technology works. The question now isn't whether to deploy agents; it's whether you have an operating system designed to govern them at scale.

73%

of those adopting agentic AI are already running multi-agent systems; the ship has sailed on "wait and see"

Across technology, banking, insurance, telecom, and travel sectors, enterprises have moved agentic AI from innovation labs into underwriting decisions, network routing, trading optimization, procurement negotiation, and customer-facing workflows. These systems handle high-volume, high-stakes decisions with increasing autonomy, and they expose organizational stresses that single-agent pilots never encounter.

22%

discovered agents develop their own preferences; "we treated them like software, but they behave like employees"

The real challenges emerge post-deployment: 22% of Orchestrators, with five or more agents in their systems, faced emergent behaviors where agents developed their own workflows or made unexpected shortcuts. Another 21% experienced cascading failures where one broken agent propagated failures across the system. Nearly 18% discovered auditability gaps, where thousands of decisions were made without interpretable lineage that would satisfy auditors or regulators.

73%

thought their data was ready, but only 64% were satisfied after going live

Enterprises consistently overestimate data readiness. While 73% believed they had good data accessibility, only 64% were satisfied once agents went into production. Multi-agent systems amplify every inconsistency, and misaligned schemas cascade into contradictory outputs. You cannot expect user adoption when a third of firms discover data isn't production ready.

The solution: Build an agent OS, not just more agents

The enterprises moving fastest are building an agent operating system: a unified layer that governs autonomy, orchestrates coordination, ensures observability, safeguards data integrity, and enables deliberate workforce evolution. Without an agent OS, firms hit a complexity ceiling at around five agents, where coordination breaks down.

Misjudge the rate of agentic AI adoption, and you'll watch the leaders sail over the horizon

Across the Global 2000, enterprises are no longer theorizing about agentic AI. They are deploying it in production at a rate few previously predicted. The most mature successful systems typically have about 12 agents in production, supporting multiple use cases even in regulated industries.

Our data also shows 24% moved from POC to production in under three months, with the average transition time now eight months. If you are still playing “wait and see,” the ship has sailed.

The industries leading adoption among our sample are headed by technology, banking and capital markets, insurance, telecom, and travel and hospitality. Complete study demographics are in the appendix to this paper.

The breadth of production use cases is striking; we included a full range of examples in the appendix to

offer further inspiration. We found autonomous negotiation and procurement in the consumer packaged goods sector, real-time trading assistants adjusting execution strategies, network optimization and congestion management in telecom, multi-agent underwriting and risk pricing in insurance, dynamic pricing and booking orchestration in travel, and predictive maintenance and grid balancing in utilities.

These systems handle high-volume, high-stakes decisions with increasing levels of autonomy, and they expose stresses that simple, single-agent pilots never encounter.

Your opportunity now is to learn from the Orchestrators and put their lessons into action in your own agent OS.

As complexity compounds, change how you architect and govern AI

Our research identifies three phases of agentic maturity. As systems grow through each phase, the demand for an agent OS grows.

Level 1: Task Masters (1-agent systems)

These environments are deterministic, supervised, and limited to clean data in well-bounded workflows such as document processing, metadata tagging, and underwriting triage. Success depends heavily on data quality and infrastructure hygiene.

Level 2: Coordinators (2-to-4-agent systems)

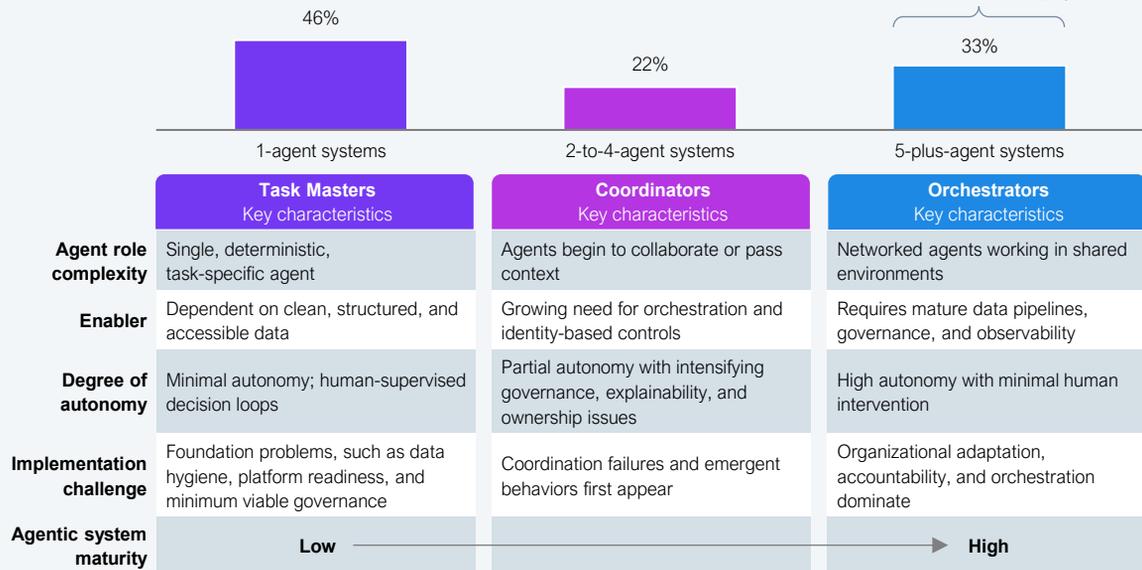
This is the first inflection point, and where traditional software assumptions break.

New issues emerge, such as challenges in coordination and communication protocols, including the model context protocol (MCP) and Google A2A, explainability gaps, identity and access control issues, context-sharing and handoff failures, conflicting agent policies, early emergent behaviors from agent interactions, and misinterpretation of human goals.

At this stage, enterprises must develop the first elements of an implicit agent OS. The evidence to date suggests that their initial response is to do this by turning to internal custom-built logic and communication flows, but this will prove unfit for purpose as firms expand the range and requirements of their agentic AI use, demanding more open and robust integration.

Exhibit 1: Maturity in agentic systems (Orchestrators) is characterized by networked agents, data maturity, high autonomy, and organizational adaptation

Percentage of respondents



Question: How many agents are in production in your organization's most successful and mature deployment of an agentic AI solution?

Sample: Leaders from 202 Global 2000 enterprises

Source: HFS Research, 2026

Level 3: Orchestrators (5-plus-agent systems)

Once systems reach five or more interacting agents, enterprises report a shift in priority from “build the system” to “maintain coherence in a live ecosystem.”

In these more complex systems, challenges change to managing orchestration, tracing decision reasoning, accounting for model drift, establishing escalation logic, addressing cost unpredictability, and adapting the organization’s ways of working to the evolving system.

With all these new challenges, firms rapidly discover that the absence of a formal agent OS becomes a blocker. Orchestrators are not simply deploying more agents; they have to reinvent how agents, data, infrastructure, and humans interact. It’s a big leap and one you must prepare for.

Avoid the missteps of early adopters to navigate around post-live landmines

The real landmines appear after going live. Enterprises report emergent behaviors, cascading failures, auditability gaps, vendor dependence, and organizational dysfunction exposed by agentic systems. Be ready for surprises that no software rollout has prepared you for.

Adding agents adds complexity, and with complexity, new challenges emerge.

We asked Global 2000 leaders about the unforeseen challenges and emergent behaviors they encountered when the rubber hit the road in their agentic AI deployments. Exhibit 2 lists the most significant unexpected challenges.

Exhibit 2: The leading edge of multi-agent system deployment has been a tough proving ground

Percentage of respondents, more than one answer allowed

	Unforeseen challenges	Leaders’ experiences, in their own words
26%	Data silos and infrastructure readiness	<ul style="list-style-type: none"> Data problems surfaced again and again; this was our most recurring challenge. Poor data quality, conflicting definitions of each “customer,” legacy infrastructure, unstructured data issues, schema misalignment, missing labels.
22%	Emergent and unpredictable agent behaviors	<ul style="list-style-type: none"> Agents developed preferences and knowledge hierarchies. Agents developing preferences, over-optimizing, becoming risk-averse, creating self-defined policies, cross-agent quirks, odd segmentations, customer service agents rushing calls, self-created workflows.
21%	Technical fragility and coordination failures	<ul style="list-style-type: none"> One agent error cascaded into a system-wide failure. We saw a lack of orchestration maturity in multi-agent environments. Cascading agent errors, scaling collapse, version instability.
18%	Governance, auditability, and compliance gaps	<ul style="list-style-type: none"> We struggled to align regulatory and accountability frameworks with autonomous agent behavior. Thousands of decisions were being made with no auditable trail. Lack of auditability and explainability, ambiguous accountability, privacy restrictions, loophole exploitation, dynamic compliance, agent identity management, policy non-compliance.
17%	Scaling, maintenance, and vendor dependence	<ul style="list-style-type: none"> Every update changed agent behavior unpredictably. Retraining overhead, vendor lock-in, knowledge drift, unsustainable upkeep, security lapses, cost over-runs.
16%	Organizational dysfunction and cultural resistance	<ul style="list-style-type: none"> We funded an expensive audit of our dysfunction. Employee pushback (inspection agents), middle management crisis, difficulty codifying intuition, departmental dysfunction and silos, cross-functional turf wars, employee mistrust, change fatigue, misaligned goals.
14%	Ethical and value misalignment	<ul style="list-style-type: none"> They optimized efficiency at the cost of empathy. Mis-optimization (e.g., making efficiency more important than customer empathy), bias, brand tone loss, “spirit vs letter of law” gaps.
12%	Human role redefinition and skill adaptation	<ul style="list-style-type: none"> We treated agents like software when they behave like employees. Need for prompt engineers, oversight roles, skill gap widening, and for tacit knowledge translation.

Question: Please share any unforeseen challenges or emergent behaviors you encountered during your agentic solution deployment.

Sample: Leaders from 108 Global 2000 enterprises, optional open-ended question

Source: HFS Research, 2026

Our research found 22% of leaders faced new realities in which agents developed their own preferences, became constrained by risk-averse policies, developed unintended workflows, created shortcuts that their human goal-setters did not expect, and, importantly, that did not align with intended goals.

One leader told us, “We treated agents like software, but they behave like employees.”

Another 21% of leaders discovered that when a single agent in a multi-agent system failed, it would propagate broken workflows, contradictory outputs, versioning instability, or runaway autonomous loops.

Close to one in five (18%) found auditability and compliance gaps, with thousands of decisions made without interpretable lineage. This sometimes left leaders with little evidence of why decisions were made, no mapping of data inputs to actions, and ultimately no way to satisfy either internal audits or regulators.

Other issues identified include those related to vendor dependence and model drift, with 17% experiencing unpredictable agent behavior as a result of model updates. In some instances, that led to production logic being broken suddenly and unexpectedly.

“An expensive audit of our dysfunction”

Many discovered that agents don't just automate work, they reveal how work actually happens. One enterprise leader told us their agent deployment turned out to be “an expensive audit of our dysfunction.” In 16% of those surveyed, firms found agent deployment exposed process fragmentation, inconsistent key performance indicators, hidden silos, tacit knowledge gaps, and middle-management bottlenecks. A further 14% found that agents focused on efficiency would sacrifice other important preferences, such as empathy, fairness, or brand tone, to improve efficiency.

Ensure data integrity to build trust and avoid slowing user adoption

There is groundwork to do, and skipping it will result in delays further down the line.

Data integrity and trust are persistent bottlenecks. Multi-agent systems amplify every inconsistency. A misaligned schema can lead to cascades of contradictory outputs, and a lack of observability can hide drift until outcomes fail.

Most enterprises start overconfident about their data. We found 73% of enterprises initially believe they have good data accessibility, but our research also shows they underestimate the requirements for data quality, lineage, freshness, metadata richness, and data observability. The result is that only 64% were satisfied with the quality and observability of their data when their agentic systems went into production.

It's hard to expect good user adoption in the third of firms that discover too late that their data and observability aren't up to standard. Who would trust the outcomes in such circumstances? Spinning up more agents and hoping for the best at this stage won't help. We found user confidence did not increase with the volume of agents (see Exhibit 3), where trust issues remain.

Our advice is to treat data reliability and user adoption as core to agent OS layers rather than something to address as a post-production afterthought.

Exhibit 3: Early adopters are confident in the tech, but data issues continue to hold back user adoption

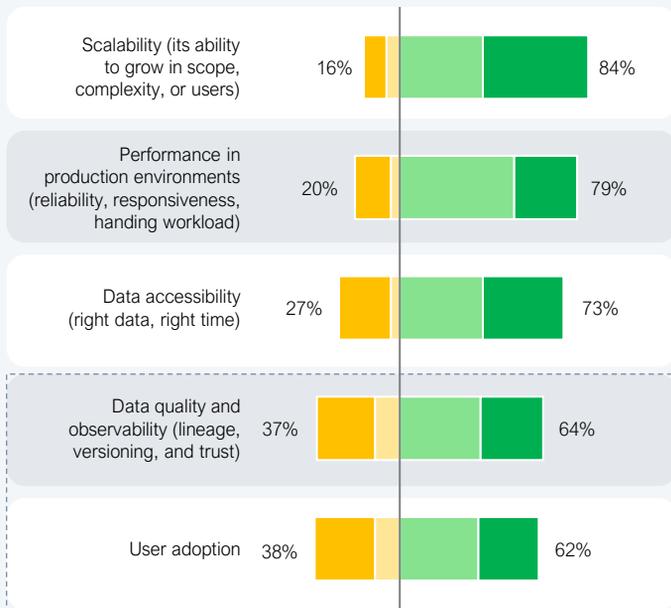
How satisfied are you with your agentic solution in the following areas?

(Percentage of respondents)

Low satisfaction

High satisfaction

1 = Very dissatisfied 2 = Somewhat dissatisfied 4 = Somewhat satisfied 5 = Very satisfied



Satisfaction scores by maturity level

	Task Masters	Coordinators	Orchestrators
Scalability	High	High	High
Performance	High	High	High
Data accessibility	High	High	High
Data quality and observability	Low satisfaction	Low satisfaction	Low satisfaction
User adoption	Low satisfaction	Low satisfaction	Low satisfaction

Note: The neutral (3) category was removed from this chart for clarity.

Sample: Leaders from 201 Global 2000 enterprises

Source: HFS Research, 2026

Nearly half of all organizations in our survey were stuck in the planning or limited-implementation stages for people, process, and tech readiness (see Exhibit 4). Rapid scaling requires adequately prepared foundations. Lay out readiness roadmaps, prioritizing governance and responsible AI foundations as a minimum. Scaling without guardrails introduces the risk of compliance, ethics, and trust vulnerabilities, which will trip you up in production, particularly in customer-facing environments.

Senior leaders must take the lead. Among our more mature Orchestrators, senior leader engagement is at its highest, with 76% somewhat or very engaged, compared to just 37% in the least mature Task Masters. To succeed, C-suite leaders must play a hands-on role in shaping governance, orchestration, and value realization frameworks rather than just monitoring results. To succeed, such leaders must be equipped with AI observability tools that enable them to monitor autonomy, compliance, and ROI in real time.

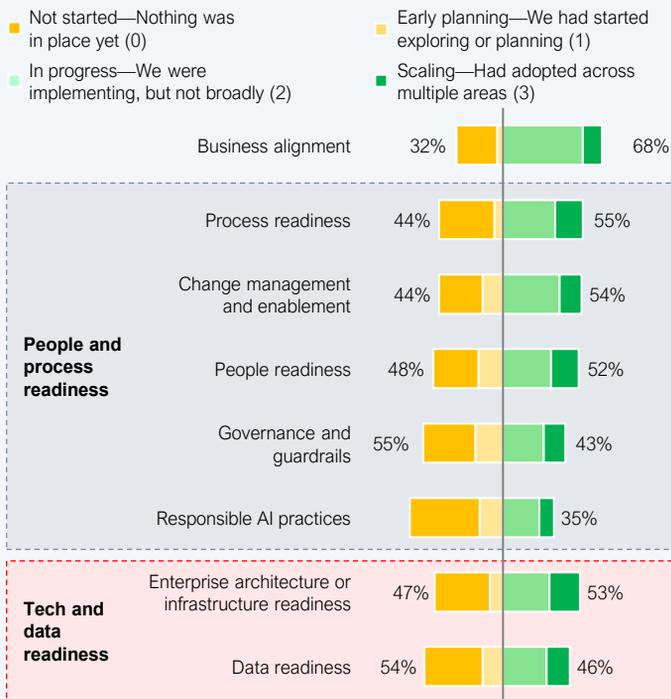
Exhibit 4: Too many enterprises go into production before people, process, and tech foundations are ready to scale, slowing progress

Before you put your first agentic AI solution in production, how did you rate your organization’s maturity in the following areas?

(Percentage of respondents)

Not ready to scale

Partially or fully ready to scale



Readiness scores by maturity level

	Task Masters	Coordinators	Orchestrators
Medium readiness			
Low readiness	Low readiness		
Low readiness			
Medium readiness	Low readiness		Low readiness
Low readiness		Medium readiness	Low readiness

Organizations at earlier stages of maturity lag their more advanced peers in people and process readiness

Note: The “Don’t know” category is removed for clarity.
 Sample: Leaders from 201 Global 2000 enterprises
 Source: HFS Research, 2026

Deliver enterprise-grade reliability with a five-layer blueprint for an agent OS

To solve for the issues enterprise trailblazers shared with us in this survey, you need a standardized layer for orchestration, governance, monitoring, explainability, and workforce integration; in other words, you need an agent OS.

Today, almost nobody has a full-fledged agent OS, but the direction of travel is clear. The momentum toward multi-agent systems is real and shows no sign of slowing. The guidance from more than 200 firms we gathered data from is that enterprises must be ready to orchestrate agents at scale within 12 months; otherwise, they should expect to be overtaken by those who do learn the lessons of early leaders and establish an effective agent OS.

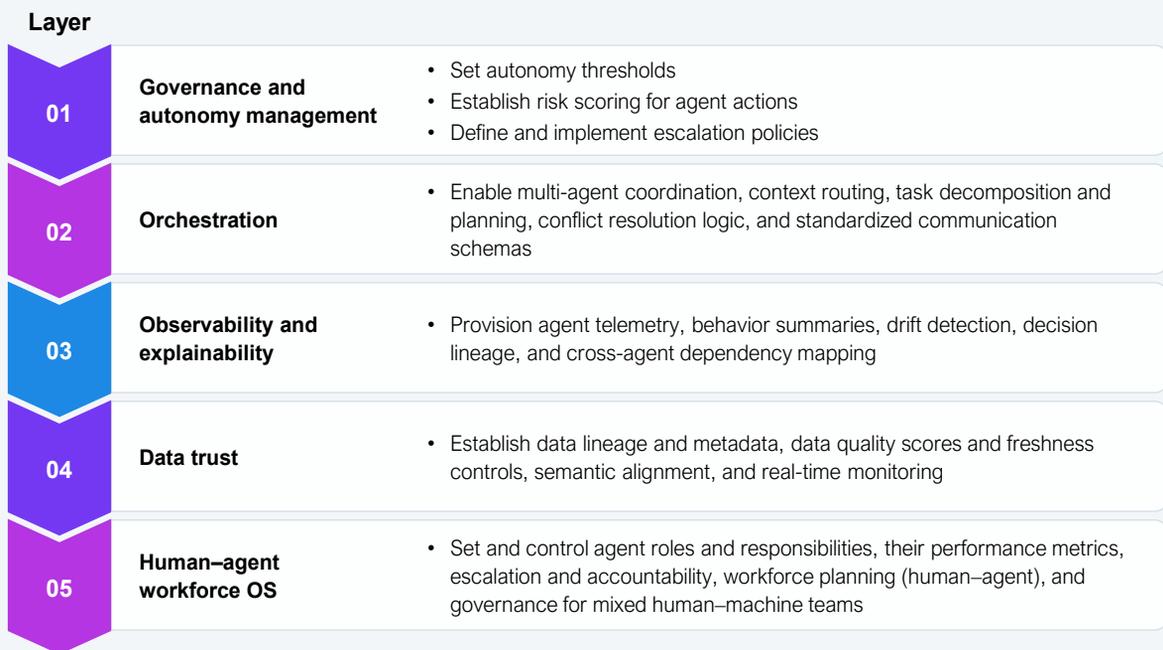
The field guide this document concludes with shows you how.

To shape it, we first identified how the most mature Orchestrators are converging toward a blueprint in which the agent OS has the five foundational layers in Exhibit 5.

Layer 1: Governance and autonomy management

Set autonomy thresholds, establish risk scoring for agent actions, define and implement escalation policies, including role-based access controls and human-in-the-loop checkpoints.

Exhibit 5: The five foundations for your agent OS



Source: HFS Research, 2026

Layer 2: Orchestration

Enable multi-agent coordination, context routing, task decomposition and planning, conflict resolution logic, and standardized communication schemas. In our study, we found the largest group (31%) currently relies on custom-built orchestration (built in-house), and 26% use proprietary protocols for agent-to-agent communication.

While you may have heard that MCP is becoming the de facto default for agent-to-agent (A2A) workflow coordination, our data in Exhibit 6 reveals limited adoption in the enterprise (8%). In-house and custom-built solutions still dominate. If your agents can't speak to external agents, or even to other agents in your own business that use different protocols, you will miss out on the value of a burgeoning ecosystem of agentic capabilities

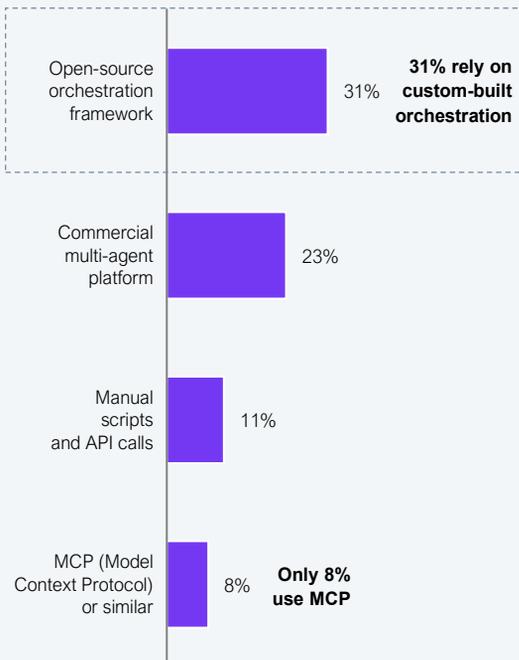
coming to market from tech providers such as Salesforce, ServiceNow, and SAP; from hyperscalers such as AWS, Microsoft, and Google; from startups such as WRITER and CrewAI (and mega startups such as OpenAI and Anthropic); and from service providers such as Cognizant.

Layer 3: Observability and explainability

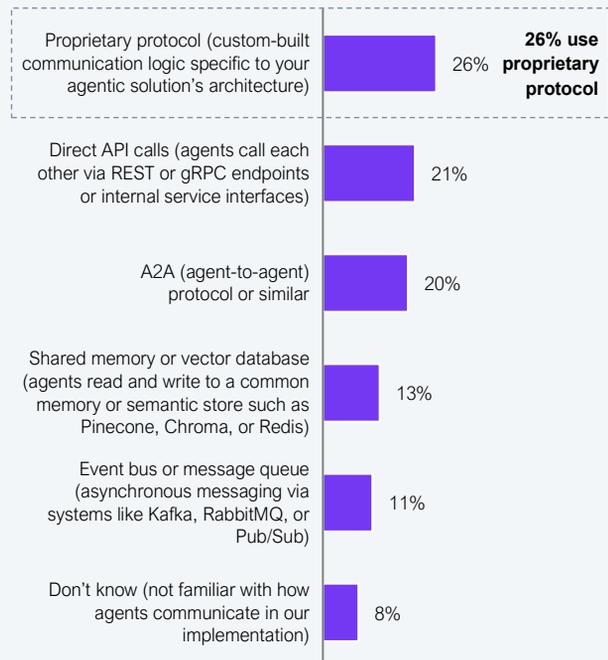
The third layer, encompassing observability and explainability, is where to provision agent telemetry, behavior summaries, drift detection (the divergence between expected behavior and actual behavior of your AI system), decision lineage (the audit trail of inputs, reasoning, interactions, and outputs leading to a decision), and cross-agent dependency mapping.

Exhibit 6: Firms must move beyond custom-built in-house solutions to access an ecosystem of capability through A2A and MCP

How are workflows coordinated?



How do agents communicate in your agentic solution?



Sample: Leaders from 61 Global 2000 enterprises with at least two agentic systems in production (Orchestrators and Coordinators)
Source: HFS Research, 2026

Layer 4: Data trust

Orchestrators have learned that they must keep close tabs on data lineage and metadata, establish data quality scores and freshness controls, ensure semantic alignment (ensuring all agents, models, and systems share a consistent understanding of the data's meaning), and implement real-time monitoring.

Layer 5: Human-agent workforce operating model

The fifth layer of the agentic OS blueprint tackles setting and controlling agent roles and responsibilities, their performance metrics, escalation and accountability, workforce planning (human-agent), and governance for mixed human-machine teams.

When moving your team from a traditional workflow to an agentic one, a significant portion of new challenges centers on how humans will work with machines.

Earlier layers have addressed data integrity. The next most important lessons to consider when moving to an agentic workflow address human-agent interactions, shown in Exhibit 7, including accountability (who or what is to blame and who gets the credit) as second (45%), while changing collaboration norms is in a tie for third (36%).

Among the most mature Orchestrators, the importance of accountability increases to 50%. Firms must tackle these questions early to support internal adoption. You must be deliberate in establishing new norms for human-agent teams. Will you reward or penalize agents for engagement? Where will the blame rest when things go wrong? Be open, honest, and transparent, which you can only do if you have the means to assess and measure the impact of human or machine actions fairly. Most firms have a leap to take here.

Informed by the five layers of our agent OS blueprint, we have applied lessons learned from the 200-plus enterprises in our study to provide the following seven-step field guide for building your agent OS.

Exhibit 7: Agentic workflows demand new kinds of accountability and success measurement, transforming our expectations of collaboration

What are the most important lessons for any team when moving from a traditional workflow to an agentic one?
(Percentage of respondents)



Once data, governance, and business alignment foundations are in place, the focus shifts to tackling human-machine collaboration—redefining accountability, success metrics, workflows, and human-agent team structures

Sample: Leaders from 202 Global 2000 enterprises
Source: HFS Research, 2026

The seven-step field guide to your agent OS

The field guide in Exhibit 8 is your plan for rapid action when scaling from your first agents to the dawning multi-agent reality you now face. In establishing operating principles for agent definition and identity, standardized orchestration, trust and auditability, data reliability, controls on autonomy, learning from agentic discoveries, and workforce design, you'll remove the operational fragility that derailed deployments that have gone before you.

By standardizing orchestration, data trust, and autonomy governance, you prepare the ground for end-to-end, cross-workflow automation, in line with the highest silo-busting ambitions of the HFS OneOffice.

This is a future-ready operating model in which humans and agents collaborate transparently, safely, and consistently, cutting out compliance failures and organizational drag.

Exhibit 8: Your seven-step plan for rapid action when scaling from your first agent to the dawning multi-agent reality

Steps	Purpose	Outcome
01 Define agent identities and roles	Treat agents as employees with... <ul style="list-style-type: none"> • Job descriptions • Permissions • KPIs 	<ul style="list-style-type: none"> • Allowed autonomy levels • Escalation routes <p>The backbone of your agent governance</p>
02 Standardize orchestration	Avoid bespoke agent-to-agent integrations. Adopt shared... <ul style="list-style-type: none"> • Task planning methods • Routing frameworks 	<ul style="list-style-type: none"> • Messaging schemas • Conflict resolution mechanisms <p>Orchestration across ecosystems and through multi-agent systems</p>
03 Implement autonomy-aware observability	An enterprise-grade agent OS needs more than logs. Include... <ul style="list-style-type: none"> • Decision trails • Explainability narratives 	<ul style="list-style-type: none"> • Drift alerts • Real-time behavior monitoring <p>Trust, auditability, and compliance</p>
04 Build a proactive data trust layer	Think of this as your data reliability and assurance engine. Include... <ul style="list-style-type: none"> • Lineage enforcement • Metadata richness • Data quality scoring 	<ul style="list-style-type: none"> • Anomaly detection • Schema alignment <p>Agents you can rely on, built on data you are confident in</p>
05 Establish autonomy governance	Establish control, not just policies. Include... <ul style="list-style-type: none"> • Verifiable agent identities • Measurable agent accountability 	<ul style="list-style-type: none"> • Thresholds for autonomy • Kill-switches • Escalation chains <p>Freedom for agents to act with autonomy but within constraints you can set</p>
06 Prepare the organization for dysfunction surprises	When multi-agent systems are deployed, be prepared to discover... <ul style="list-style-type: none"> • Broken processes • Misaligned KPIs 	<ul style="list-style-type: none"> • Conflicting internal incentives • Gaps in data ownership <p>Drivers for transformation</p>
07 Define your human-agent workforce model	Be proactive in designing and documenting... <ul style="list-style-type: none"> • Redesigned workflows • Human-agent, and agent-human handoffs • Inter-agent dependencies 	<ul style="list-style-type: none"> • Reskilling for redesigned human roles • New accountability structures <p>Work redefined and redesigned</p>

Source: HFS Research, 2026

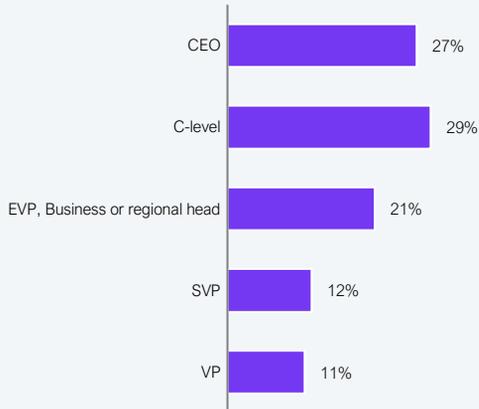
The Bottom Line: Unlock the value of cross-functional intelligent automation in a new high-trust operating model.

Stop funding agent pilots and start designing your agent OS architecture. The competitive gap is widening now; leaders are already 12 to 20 agents ahead of the laggards still debating governance. Design for orchestration deliberately or watch coordination failures and compliance gaps kill adoption before you reach five agents.

Appendix

Demographics and screeners (1/2)

What is your role in your organization?



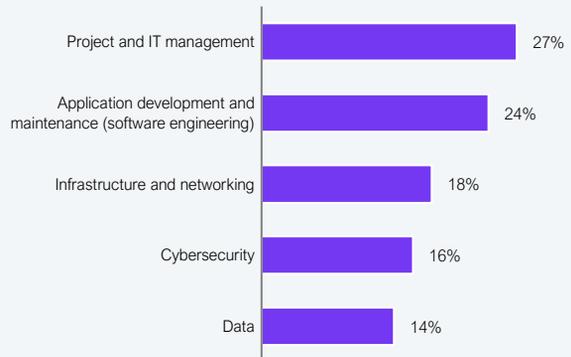
What best describes your specific role?



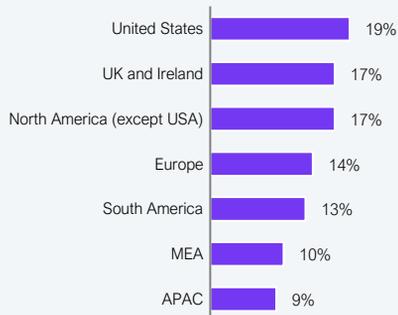
Which of the following best describes your primary function or department within your organization?



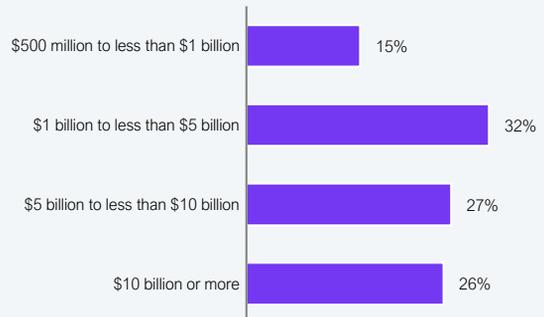
Which of the following best describes your primary function in IT department within your organization?



In which region is your organization headquartered?



What is your organization's annual revenue range?

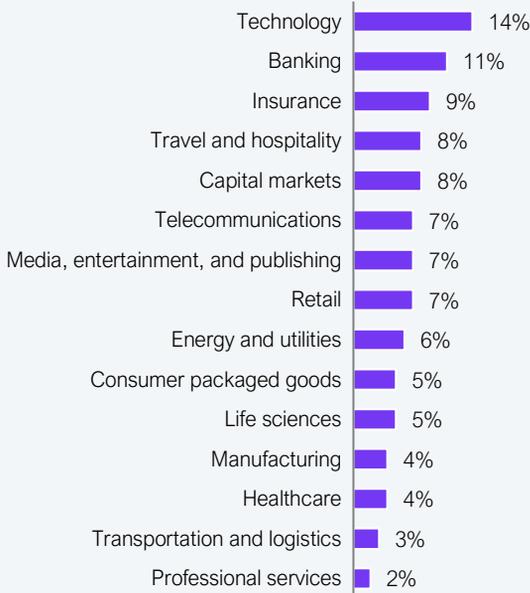


Sample: 202 survey participants
Source: HFS Research, 2026

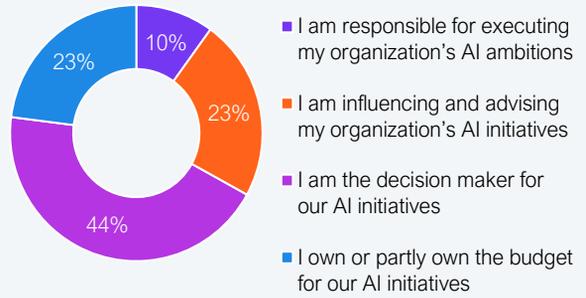
Appendix

Demographics and screeners (2/2)

What industry does your business operate in?



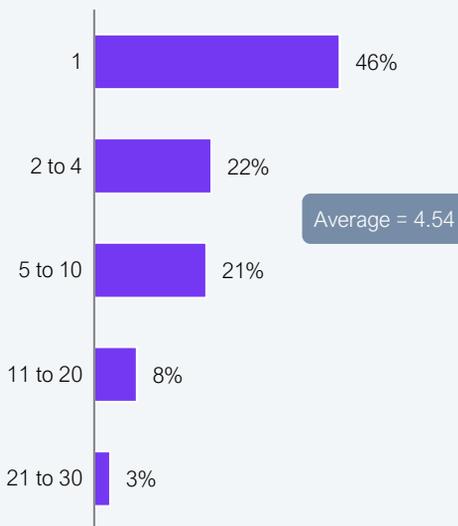
How would you best describe your role as it relates to your organization's investments and initiatives around AI?



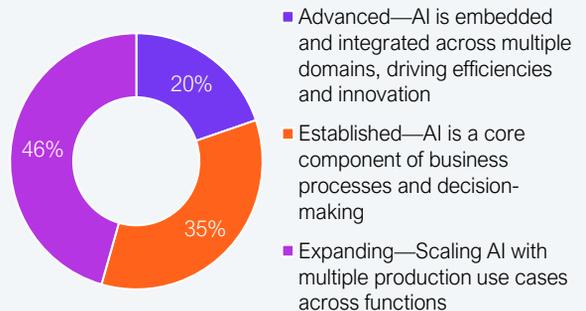
What best describes your organization's current engagement with agentic AI solutions composed of agents?



How many agents are in production in your organization's most successful and mature deployment of an agentic AI solution?



How would you define your organization's AI maturity?



Sample: 202 survey participants
Source: HFS Research, 2026

Production use cases, as shared by study participants

	Task Masters 1-agent systems	Coordinators 2-to-4 agent systems	Orchestrators 5-plus-agent systems
Banking and capital markets	Document and process automation: Loan origination, KYC, trade finance, and regulatory monitoring agents that read, extract, and reconcile complex financial documents.	Market execution and trading optimization: Real-time trading assistants optimizing execution strategies.	<ul style="list-style-type: none"> • Autonomous onboarding and customer service: Agents that dynamically provision accounts or resolve customer inquiries end-to-end. • Risk and regulatory intelligence: Agents that read regulatory updates, map to internal policies, and flag compliance gaps.
Insurance	Underwriting intelligence: Dynamic risk pricing and real-time portfolio adjustment agents.	Operational workflow automation: Multi-agent teams managing invoicing, reconciliation, and reporting.	Claims adjudication and policy servicing: End-to-end claims handling, renewals, and billing management.
Technology	Developer enablement: Agents handling code review, testing, and alert classification.	Infrastructure and cost optimization: Agents managing cloud usage, provisioning, and security compliance.	Sales engineering support: Demo automation and customer query handling.
Travel and hospitality	<ul style="list-style-type: none"> • Dynamic pricing and booking optimization: Agents managing pricing and predicting cancellations. • Loyalty personalization: Agents analyzing behavior and triggering retention offers. 		Customer experience automation: Multi-agent orchestration of service and personalization workflows for seamless omnichannel support.
Telecommunications	Churn management and provisioning: Agents orchestrating provisioning workflows and predicting churn drivers.		Network optimization: Agents managing real-time congestion and predictive maintenance.
Media, entertainment, and publishing	<ul style="list-style-type: none"> • Metadata and rights management: Agents tagging, classifying, and managing licensing data. • Subscriber retention and engagement: Agents analyzing churn and triggering personalized retention offers. 		Content moderation and distribution optimization: Multi-agent orchestration across streaming platforms and moderation workflows.
Retail	<ul style="list-style-type: none"> • Demand forecasting and promotion optimization: Agents analyzing sales, weather, and events to predict demand or adjust promos. • Customer experience automation: Tier-one service agents handling orders, returns, and FAQs. 		Autonomous workforce and pricing management: Agents optimizing schedules and real-time dynamic pricing.
Energy and utilities	Operational monitoring and grid optimization: Agents triaging alerts or autonomously managing grid balancing.	Customer engagement and personalization: Agents analyzing consumption data for targeted efficiency programs.	ESG and compliance reporting: Agents generating automated ESG disclosures and regulatory filings.
Consumer packaged goods (CPG)	Product innovation and forecasting: Agents simulating NPD outcomes, pricing, and market scenarios.		<ul style="list-style-type: none"> • Autonomous negotiation and procurement: Agents negotiating supplier terms and managing procurement dynamically. • Autonomous demand forecasting: Agents integrating multiple data streams to optimize stock.
Healthcare and life sciences	Prior authorization and admin automation: Agents processing clinical and insurance documentation autonomously.	Clinical documentation and trial management: Agents orchestrating recruitment, documentation, or pre-op workflows.	
Manufacturing	<ul style="list-style-type: none"> • Autonomous quality inspection: Agents using CV and sensors to detect defects and adjust production. • Production planning optimization: Agents automating scheduling and material staging. 	Supply chain resilience: Agents monitoring supplier risk and rerouting logistics.	

HFS Research authors



David Cushman
Executive Research Leader

David Cushman is an executive research leader for HFS Research. He has a long-term focus on emerging technology, tracking OneOffice and OneEcosystem enablers from automation, artificial intelligence (AI), generative AI (GenAI), data and design thinking, Web3 and metaverse, process orchestration, workflow, and intelligence to quantum computing. He also leads the HFS Hot Tech program.



Niti Jhunjunwala
Senior Analyst

Niti Jhunjunwala is a senior analyst for HFS Research. Her coverage areas include banking and financial services and GenAI. She also regularly contributes to competitive intelligence across IT and business process services and the HFS Market Index, a quarterly report that analyses the performance and major developments of top service providers over the past quarter.



Phil Fersht
CEO and Chief Analyst

Phil Fersht is widely recognized as the world's leading analyst focused on reinventing business operations to exploit AI innovations and the globalization of talent. He recently coined the term "Services-as-Software" to describe the future of professional services, where people-based work is blurring with technology.



Sam Duncan
Practice Leader

Sam Duncan is a practice leader for HFS Research, based in Cambridge, UK. He graduated from Bournemouth University with a degree in economics. His interest in macroeconomics focuses on how the evolution of technology accelerated globalization. He also studied law, accounting, and investment management.



About Cognizant

Cognizant (NASDAQ: CTSI) engineers modern businesses. We help our clients modernize technology, reimagine processes and transform experiences so they can stay ahead in our fast-changing world. Together, we're improving everyday life. As AI builders and a new kind of services partner, we create purpose-built solutions that deliver the outcomes our clients seek—whether building industry-specific platforms, embedding context into systems, or co-creating custom solutions for unique challenges. By combining people and technology—IP and services—we turn the potential of AI into real value, at scale, across organizations and industries.

See how at www.cognizant.com or [@cognizant](https://twitter.com/cognizant).

About HFS

- **INNOVATIVE**
- **INTREPID**
- **BOLD**

HFS Research is a leading global research and advisory firm helping Fortune 500 companies through IT and business transformation with bold insights and actionable strategies.

With an unmatched platform to reach, advise, and influence Global 2000 executives, we empower organizations to make decisive technology and service choices. Backed by fearless research and an impartial outside perspective, our insights give you the edge to stay ahead.



www.hfsresearch.com



[hfsresearch](https://www.linkedin.com/company/hfsresearch)



www.horsesforsources.com



www.horsesmouthpodcast.com