

## MARKET IMPACT REPORT

# Stop scaling agentic AI on operating models built for control

How to join the 6% scaling enterprise-wide  
scope, not just agent activity

### Authors:

Dana Daher, Executive Research Leader

Phil Fersht, CEO and Chief Analyst

Hridika Biswas, Associate Practice Leader

# Contents

	Page
Executive summary	03
Most enterprises are scaling agent activity, not autonomy	05
Strategy is outpacing structure, and autonomy breaks where ownership is unclear	11
The foundation is too fragile for autonomy	15
The real barrier to agentic is not code; it's culture	19
Autonomy will be built in workflows, not workarounds	25
The operating model reset: How to join the 6%	27
The Bottom Line	28
HFS Research authors	29
Participant demographics	30
About Infosys	32
About HFS	33

## Executive summary

Agentic AI marks a transition from tools that assist work to systems that participate in it. In 2026, enterprises are no longer experimenting at the edges; they are embedding intelligence into execution itself. The question is no longer whether agents can act, but whether organizations are prepared for what it means when they do.

This HFS study, conducted in partnership with Infosys and informed by more than 500 Global 2000 enterprise decision makers, examines how enterprises are deploying agentic AI, where it is delivering value, what is constraining autonomy, and what separates isolated pilots from enterprise-wide scale.

The findings show that progress is real but uneven[HS1.1][HB1.2]. Adoption is accelerating, yet autonomy remains constrained by how work is structured, governed, and trusted. Most enterprises are building agent capability faster than they are reshaping the systems those agents must operate within.

## Key findings

1

### Autonomy is the headline, but scope is still the constraint.

Eighty percent (80%) of enterprises remain in the Exploring or Emerging stages, only 14% have reached Scaling, and just 6% qualify as Pioneers. Even now, only 16% report enterprise-wide scope, which keeps autonomy contained.

2

### Most "agentic" still looks like controlled automation.

Sixty percent (60%) say their most advanced agents are performing simple, rules-based tasks, and only 16% report enterprise-wide deployment. Agents are present, but they are largely confined to safe, scripted terrain.

3

### Ambition is enterprise-grade, but ownership is still fragmented.

Almost half (47%) are pursuing agentic AI to rethink business models or operating structures, yet 28% report competing or fragmented mandates, and 20% are still debating who owns outcomes.

4

### Foundations are the ceiling; data readiness and permissions are the choke points.

Forty-four percent (44%) cite data and infrastructure gaps, only 16% report real-time data availability, and just 12% are comfortable granting agents broad access to sensitive data.

5

## Culture is not rejecting agents, but it is negotiating boundaries and visibility.

Fifty-two percent (52%) report resistance to integrating agents into workflows, 40% have no formal training or guidelines for human-agent collaboration, and 80% agree agents need access to detailed performance and activity data to optimize work.

6

## Workflow redesign is the real scaling lever, not adding more agents.

Thirty-seven percent (37%) are layering agents onto existing workflows, and 32% are making limited optimizations, while only one in five are redesigning workflows for autonomy. The strongest accelerators are process redesign (35%) and orchestration roles (30%).

Agentic AI is no longer a feature upgrade. It is a structural change in how decisions are made, how work flows, and how responsibility is assigned between humans and machines. The enterprises that navigate this shift well will do so by designing for coordination, trust, and adaptability, not by adding more intelligence to systems built for control.

## Most enterprises are scaling agent activity, not autonomy

For enterprise leaders scaling AI across operations and governance, Agentic AI adoption is accelerating, and deployment is expanding across tasks, teams, and workflows. But a familiar pattern is emerging with implementation moving faster than enterprise redesign. Agents are being embedded into execution, yet decision rights, accountability, and governance remain largely unchanged. Without structural shifts in how authority and ownership are defined, organizations are increasing activity within existing models rather than enabling true autonomy.

### Belief is widespread; maturity is not


Agentic AI has already won the belief battle. Ninety-two percent (92%) of executives say it will

reshape the fundamentals of how work gets done (see Exhibit 1). The momentum is real, but the next question is the one most enterprises are now running into: What does “ready” actually look like when agents move from demos into day-to-day execution?

That is where agentic maturity matters. It is not a measure of how many agents an enterprise has deployed. It is a measure of whether agents can expand beyond isolated tasks into real workflows, with clear ownership, governed behavior, trusted data access, and defined human intervention when things break. When those conditions are missing, adoption still happens, but it stays contained.

### Exhibit 1: Agentic AI has won belief; now it has to earn readiness

To what extent do you agree with the following statements?



**Agentic AI will reshape the fundamentals of how work gets done in our organization**

Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

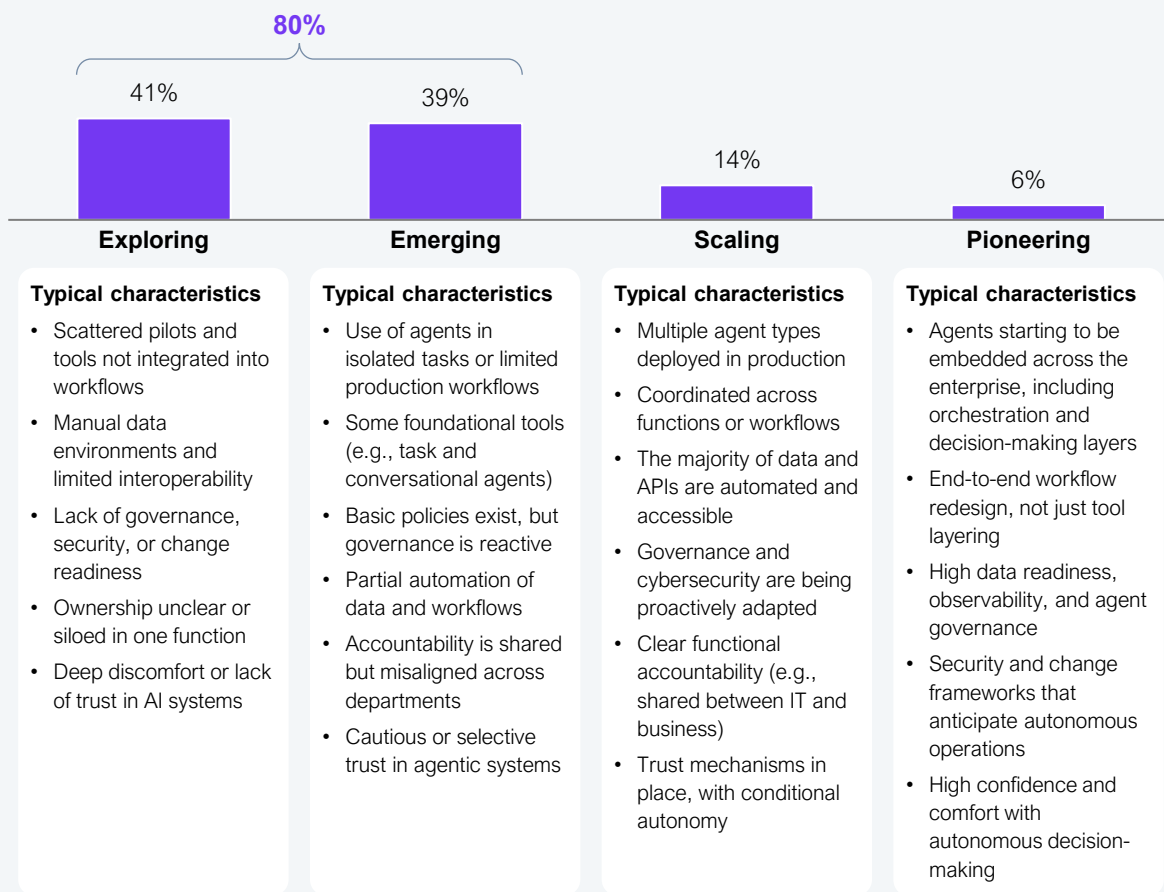
The maturity distribution makes that containment visible. Enterprises cluster into four stages: Exploring (41%), Emerging (39%), Scaling (14%), and Pioneering (6%). Eighty percent (80%) remain in the first two stages, where progress is still concentrated in bounded use cases and value is measured primarily through efficiency gains rather than enterprise-wide execution change (see Exhibit 2).

This maturity did not appear overnight. The pioneers are also the enterprises that were

already ahead on GenAI maturity. They are more than 20 times more likely to have GenAI embedded across core workflows and decision making than organizations still exploring, consistent with last year's [HFS-Infosys AI maturity research](#). The message is simple: agentic progress compounds. Enterprises that never moved GenAI beyond pilots are now discovering that scaling agents is easy, but scaling workflows, ownership, and trust is the real work.

## Exhibit 2: 80% of organizations are in the early phases of maturity

### Organizational Agentic AI maturity



Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

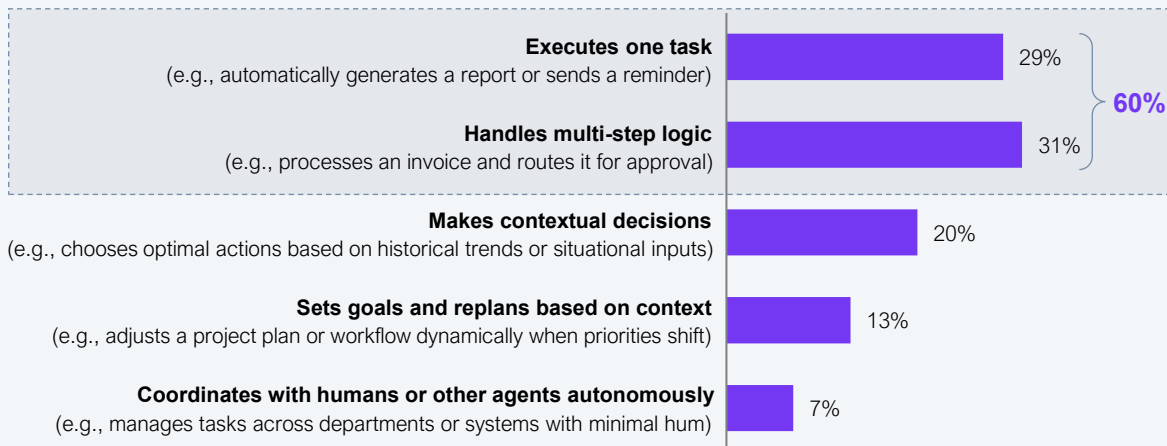
## Most "agentic" today is still controlled delegation

Enterprises are choosing reliability over reach, and the data makes that choice visible. Sixty percent (60%) say their most advanced agents are still rules-based and focused on bounded tasks like summarization, ticket handling, and transaction processing (see Exhibit 3). That is not a weak starting point; it is a predictable one. It keeps risk low, value measurable, and failure containable.

But it also exposes what maturity really demands. Task automation is not autonomy; it is permissioned execution. The maturity jump shows up when agents move from executing instructions to handling multi-step logic, making contextual decisions, and coordinating across humans and systems. These capabilities concentrate in the Scaling and Pioneering cohorts (see Exhibit 3). The message is simple: Autonomy does not arrive when agents get smarter. It arrives when enterprises change what agents are allowed to do and build the governance, data access, and accountability to let them do it.

### Exhibit 3: Most organizations are deploying simple, rules-based agents focused on basic tasks

#### What best describes the current capability of your most advanced AI agent?



	Exploring	Emerging	Scaling	Pioneering
Executes one task	41%	27%	11%	0%
Handles multi-step logic	38%	31%	19%	10%
Makes contextual decisions	12%	26%	28%	17%
Sets goals and replans based on context	6%	10%	32%	45%
Coordinates with humans or other agents autonomously	3%	7%	10%	28%

Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

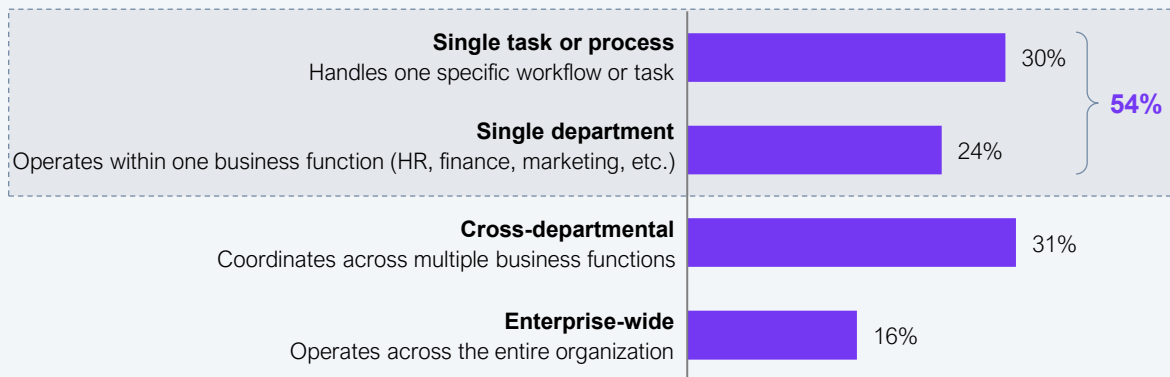
## Scope is the maturity test, and most deployments are still contained

If you want one clean maturity signal, it is scope. Only 16% report enterprise-wide deployment today (see Exhibit 4). Most deployments remain confined to tasks, processes, or single functions, which limits coordination and keeps value localized.

This is not a model capability problem. It is a delegation problem. Expanding scope forces decisions about decision rights, escalation, auditability, and shared accountability. Enterprises that cannot answer those questions will keep agents contained, even when adoption looks broad.

### Exhibit 4: More than half of agents remain limited to single departments or workflows

#### What is the organizational scope of your most advanced AI agent?



#### Scope of agentic by maturity

Exploring	Emerging	Scaling	Pioneering
Single task or process <b>(41%)</b>	Cross-departmental <b>(36%)</b>	Cross-departmental <b>(44%)</b>	Enterprise-wide <b>(48%)</b>
Single department <b>(32%)</b>	Single task or process <b>(28%)</b>	Enterprise-wide <b>(31%)</b>	Cross-departmental <b>(38%)</b>
Cross-departmental <b>(20%)</b>	Single department <b>(21%)</b>	Single department <b>(15%)</b>	Single task or process <b>(10%)</b>
Enterprise-wide <b>(6%)</b>	Enterprise-wide <b>(15%)</b>	Single task or process <b>(10%)</b>	Single department <b>(3%)</b>

Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

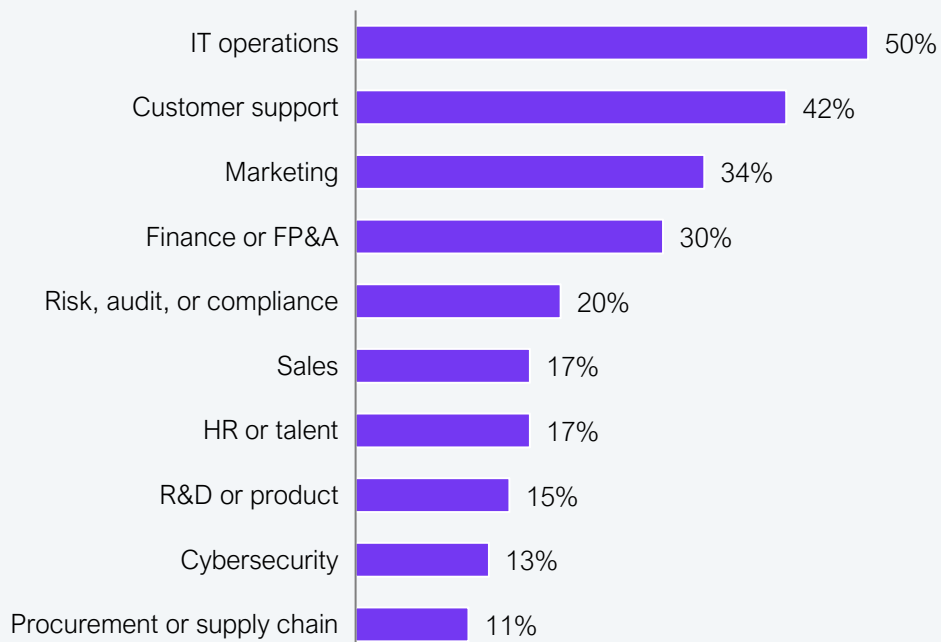
## Adoption accelerates where structure already exists

Agentic AI gains traction fastest where workflows are already structured, data is abundant, and outcomes are measurable. IT operations and customer support lead adoption, cited by 50% and 42% of enterprises, respectively (see Exhibit 5). These are domains where escalation paths are clear, risk is bounded, and governance can be operationalized.

The drop-off in areas like risk management, product development, and supply chain orchestration is equally instructive. These environments are exception-rich, judgment-heavy, and cross-functional by nature, which makes it harder to formalize ownership and decision rights. The constraint is not that agents cannot add value there. It is that most enterprises have not yet redesigned execution logic for work where trade-offs dominate and accountability is shared.

### Exhibit 5: Agentic AI success is concentrated in IT operations and customer-facing functions

In which domains has agentic AI had the most measurable success to date?



Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

## Sector-specific adoption reveals distinct maturity paths

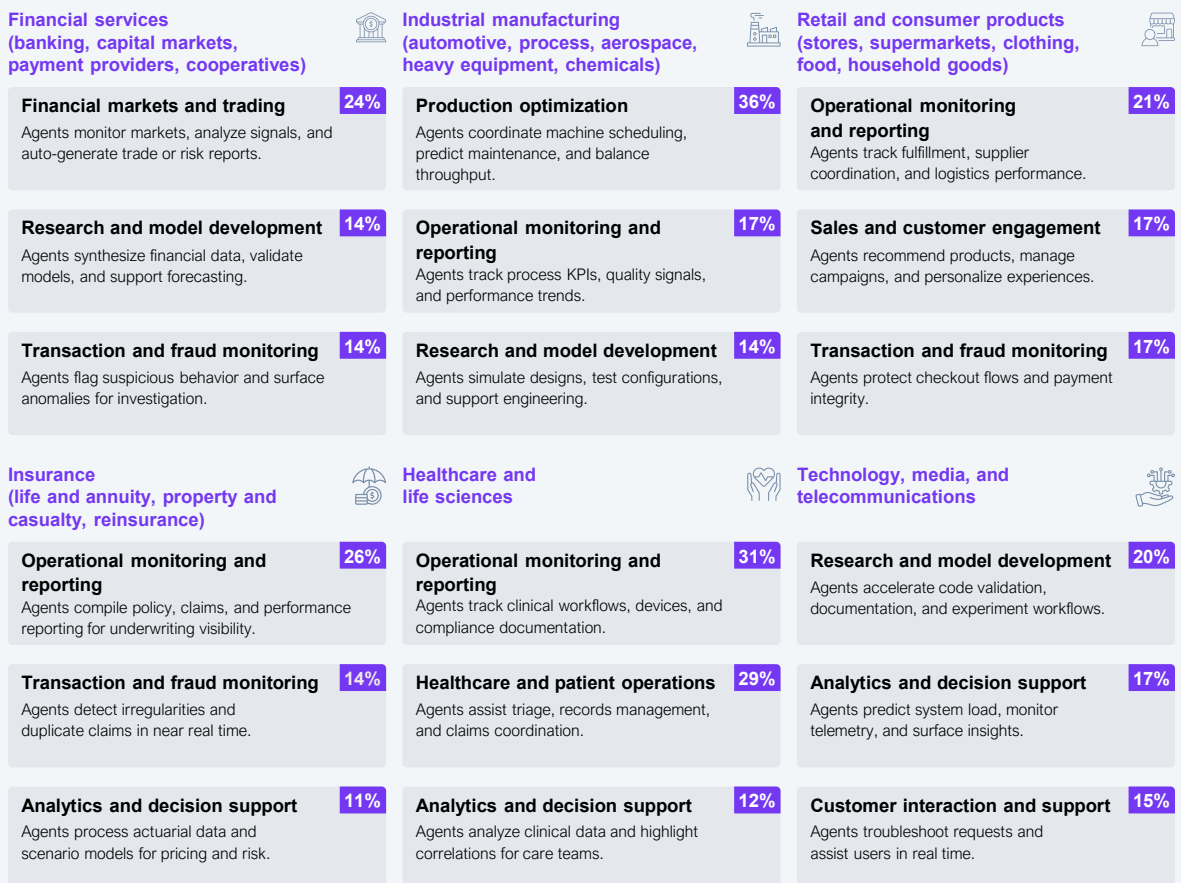
Industries are not converging on agentic AI at the same pace because they do not share the same risk thresholds. Financial services skews toward risk and compliance, healthcare and manufacturing emphasize safety and reliability, retail leans into customer engagement, and technology firms tend to experiment more broadly. Across sectors, adoption is still concentrated in bounded use cases like research, monitoring, and analytics, work that

improves speed and insight without requiring broad delegated authority (see Exhibit 6).

## Belief will not build autonomy

Most enterprises are proving that agents can work within clear boundaries. Far fewer are proving that the enterprise can work with agents across workflows, functions, and accountability lines. The maturity test is not adoption volume; it is whether agentic systems can expand scope safely, with defined decision rights, real-time data access, auditability, and escalation rules that make accountability explicit.

### Exhibit 6: Across industries, agentic adoption is led by research, monitoring, and analytics use cases



Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

## Strategy is outpacing structure, and autonomy breaks where ownership is unclear

Enterprises are aiming high with agentic AI. Many are explicitly framing it as a way to rethink how the business operates, not just how tasks get done. But the data shows a growing disconnect between strategic intent and structural readiness. Organizations are running toward autonomy without agreeing on who is steering.

### Ambition is transformational; orchestration is not

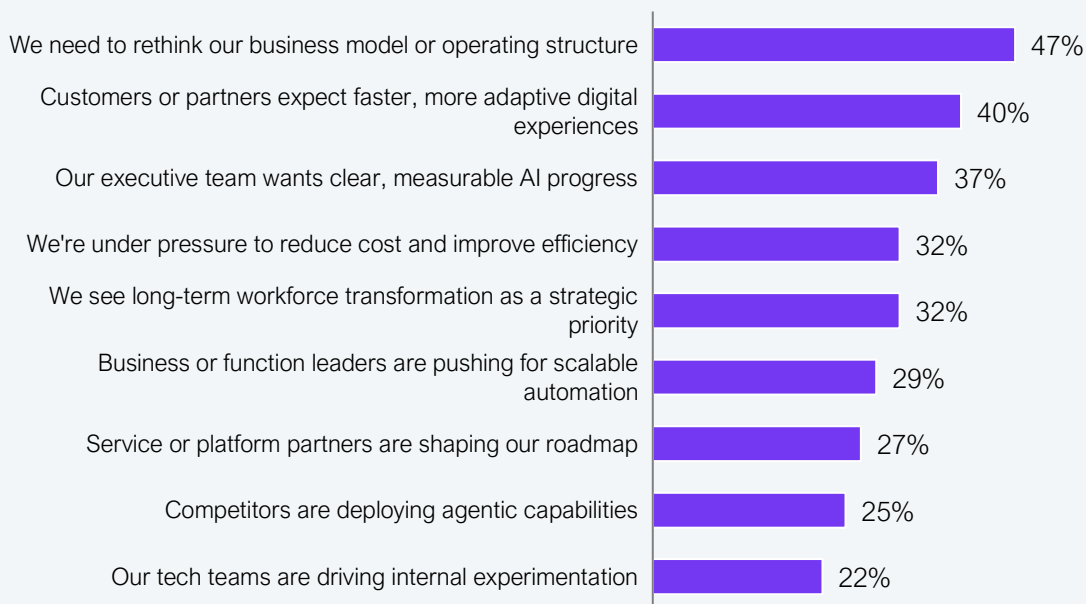
The intent behind agentic AI is not incremental. Forty-seven percent (47%) say they are pursuing it to rethink business models or operating

structures, and another 40% are aiming to meet rising expectations for adaptive digital experiences (see Exhibit 7). These are not efficiency plays. They signal recognition that competitiveness increasingly depends on enterprises that can sense, decide, and act dynamically.

Yet the operating reality does not match the ambition. Many organizations are attempting enterprise reinvention without a coordinating layer that aligns agents, humans, and accountability across workflows. Strategy is being set at the top, while execution authority remains fragmented at the edges.

### Exhibit 7: Top reasons for pursuing agentic AI center on business model reinvention

#### What are your organization's top reasons for pursuing agentic AI?



Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

## Early value shows up in efficiency, not in decision advantage

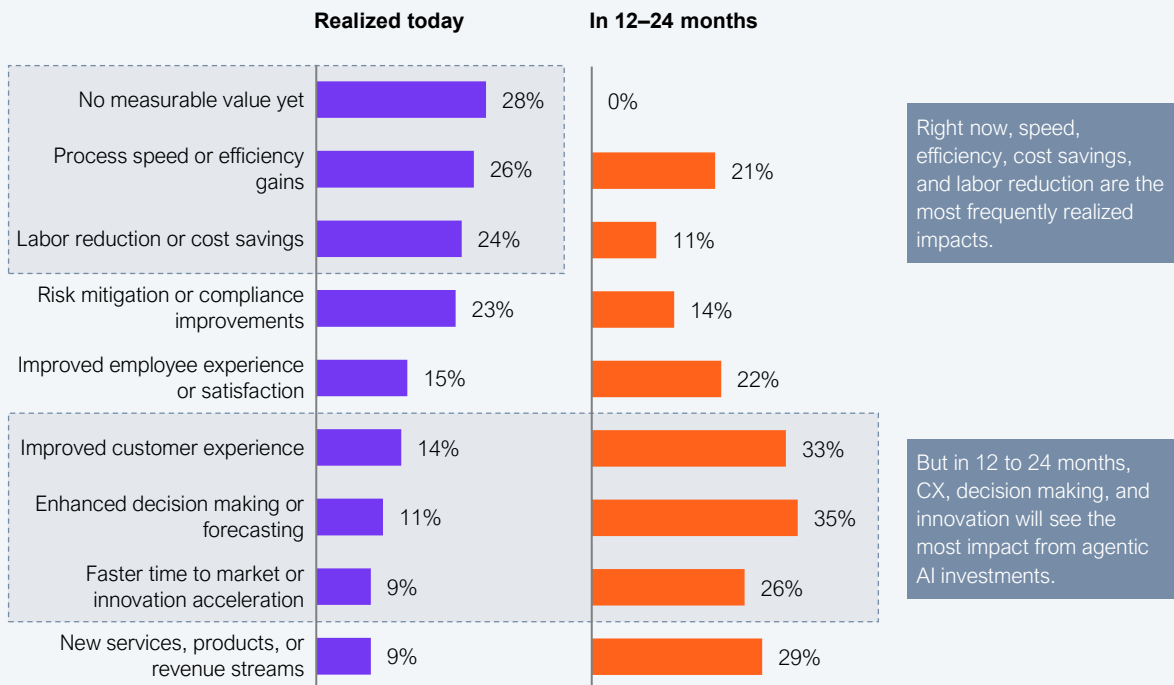
The outcomes realized so far reflect this imbalance. Enterprises report gains in process efficiency (26%) and cost reduction (24%), but far fewer see impact in decision quality, innovation, or customer experience, the areas where agentic systems should ultimately matter most (see Exhibit 8). Nearly one in three reports no measurable value yet.

This does not mean agentic AI is underperforming. It means it is being applied where a structure already exists. Efficiency improves first because it fits inside existing accountability models. Decision advantage lags because it requires new ones.

Leaders expect this to change over the next 12 to 24 months, with greater impact anticipated in forecasting, decision intelligence, and innovation. That expectation assumes a structural shift, not just better models.

### Exhibit 8: Enterprises are realizing tactical gains today but expect strategic impact tomorrow

**What business impact has your organization realized (or expects to realize) from agentic AI investments?**



Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

## Ownership ambiguity is the scaling bottleneck

Autonomy cannot scale without clarity on who owns outcomes. Yet ownership remains unresolved across much of the enterprise. Twenty-eight percent (28%) report that multiple functions are involved with fragmented or competing mandates, and another 20% say ownership is still being debated (see Exhibit 9). Only a minority reports clear, centralized accountability.

The maturity signal here is stark. In early-stage organizations, ownership is diffuse or absent. In pioneering ones, cross-functional ownership models, often through dedicated centers of excellence, are the norm. This mirrors earlier digital transformations, but with higher stakes. Agentic systems do not just execute faster; they act independently within defined bounds. Ambiguity at this layer becomes risk, not flexibility.

### Exhibit 9: Ownership of agentic AI remains fragmented across functions

#### How would you describe ownership and accountability for the agentic AI agenda in your organization today?



	Exploring	Emerging	Scaling	Pioneering
A single function owns it (e.g., IT, business, or transformation lead) with clear accountability	28%	1%	32%	3%
Multiple functions co-own it (e.g., business, tech, HR) with defined collaboration structures	19%	19%	15%	10%
Multiple functions are involved but accountability is fragmented or competing	36%	30%	8%	3%
We've formed a dedicated cross-functional task force or center of excellence	10%	0%	44%	83%
Ownership is still being debated or evolving	4%	44%	0%	0%
Ownership is still being debated or evolving	3%	6%	0%	0%

Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

## When teams cannot agree on what "agentic" means, autonomy stalls

Structural misalignment is reinforced by conceptual fragmentation. Just over half of leaders say their teams share a common understanding of what agentic AI is and how it should be applied, while nearly half remain uncertain or divided (see Exhibit 10). Without a shared definition, governance becomes reactive and trust remains fragile.

Autonomy means different things to different groups. To some, it is automation with reasoning.

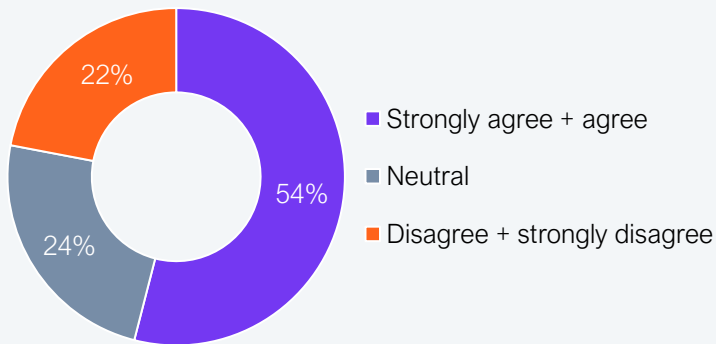
To others, it is delegated decision making. Until enterprises align on where human judgment ends and agentic authority begins, they will continue to design for containment rather than scale.

## Ambition without orchestration creates entropy

Agentic transformation cannot be driven from strategy decks or isolated innovation teams. The gap is no longer intent. It is integration. Leadership conviction is high, but ownership is fragmented, definitions are inconsistent, and decision rights are unclear.

### Exhibit 10: Teams don't have a shared understanding of what agentic is

To what extent do you agree with the following statement?



**Teams have a shared understanding of what agentic AI is and how it should be applied**

Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

# The foundation is too fragile for autonomy

Agentic AI exposes what enterprises have quietly worked around for years. Intelligence can reason, plan, and adapt, but it cannot operate reliably on data that arrives late, systems that do not interoperate, or controls designed for static workflows. As autonomy expands, foundational weaknesses stop being a technical inconvenience and become an execution risk.

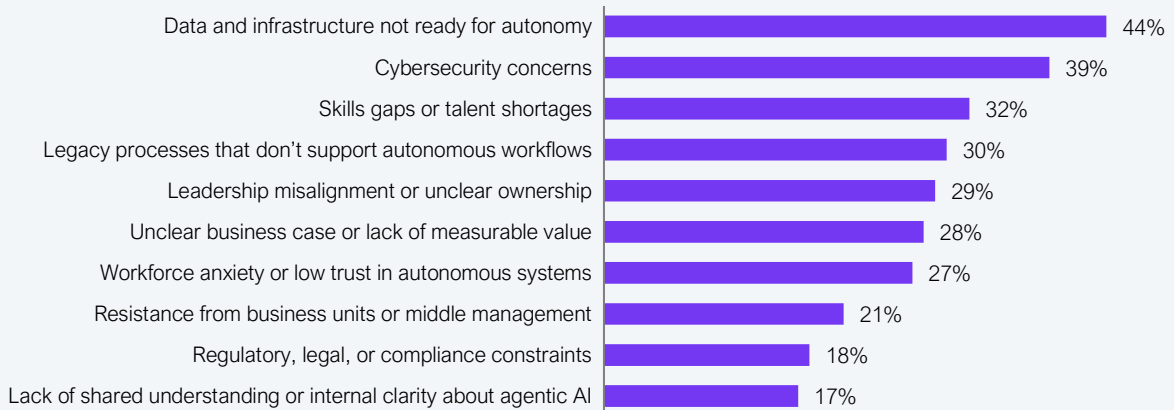
This is not a tooling problem. It is an operating environment problem. The data shows that enterprises are pushing agentic systems toward greater responsibility before the foundations needed to support delegated action are in place.

## Data and infrastructure are the hard limits on delegation

Autonomy depends on timely, contextual access to data. Yet 44% of enterprises say their data and infrastructure are not ready to support agentic AI at scale (see Exhibit 11). Only a small minority report real-time data availability across workflows, which sharply limits how much authority agents can safely be given.

**Exhibit 11: Barriers shift from technical readiness to organizational orchestration as maturity increases**

### What are the biggest barriers your organization faces in scaling agentic AI?



Rank	Exploring	Emerging	Scaling	Pioneering
1	Data and infrastructure not ready for autonomy	Data and infrastructure not ready for autonomy	Data and infrastructure not ready for autonomy	Cybersecurity concerns
2	Cybersecurity concerns	Cybersecurity concerns	Cybersecurity concerns	Regulatory, legal, or compliance constraints
3	Skills gaps or talent shortages	Skills gaps or talent shortages	Leadership misalignment or unclear ownership	Data and infrastructure not ready for autonomy

Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

This explains why so many deployments remain assistive rather than autonomous. When data is fragmented, delayed, or siloed, agents can recommend and execute tasks, but they cannot own decisions. Infrastructure reinforces this constraint. Many environments are still optimized for batch processing, human checkpoints, and sequential handoffs, not continuous execution.

## Data is not just scarce; it is selectively trusted


Even where data exists, trust remains tightly rationed. Only 53% allow agents access to

sensitive data only in predefined use cases, while just 12% are comfortable granting broad access (see Exhibit 12).

This is guarded autonomy. Agents may be capable of reasoning across contexts, but permissions keep them operating inside narrow corridors. The maturity signal is consistent: early-stage organizations are still connecting and cleaning data, scaling organizations are adding controls, and pioneers are embedding governance directly into data and execution architectures.

### Exhibit 12: Only 12% fully trust agents with sensitive data, while 53% take a highly selective approach

#### How comfortable is your organization with allowing agents to access and act on sensitive data?



	Exploring	Emerging	Scaling	Pioneering
<b>Very uncomfortable</b> We actively restrict agent access	6%	13%	6%	0%
<b>Cautious</b> Only limited access with strict oversight	23%	35%	14%	7%
<b>Selective</b> Agents access data based on clear rules and use cases	<b>63%</b>	<b>50%</b>	<b>43%</b>	34%
<b>Comfortable</b> Agents are trusted to access and use sensitive data responsibly	7%	2%	38%	<b>59%</b>

Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

## Cybersecurity models were not built for systems that decide


Agentic AI introduces a different risk profile than traditional automation. Only 16% of enterprises say their cybersecurity frameworks are ready to address agent-specific threats, while more than half acknowledge that a major redesign is required (see Exhibit 13). Static security models

assume predictable behavior. Agentic systems operate through dynamic decision chains that cross systems and adapt in real time.

Most enterprises are responding by layering new controls onto old frameworks. That approach maintains compliance, but it does not create resilience. As autonomy grows, security must evolve from perimeter protection to continuous monitoring, identity-aware execution, and behavior-based controls.

### Exhibit 13: Few enterprises have cybersecurity models ready for agentic risks

#### Which best describes your cybersecurity readiness for agentic AI?



Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

## Integration friction is slowing scale more than model capability

Even when data and security progress, orchestration remains a bottleneck. Sixty-four percent (64%) of enterprises rely on pre-built connectors or manual configuration to integrate agents, and only 13% report seamless integration with core enterprise systems (see Exhibit 14). Every brittle integration limits coordination, increases latency, and constrains cross-workflow execution.

What emerges is architectural fragmentation that mirrors organizational fragmentation. Enterprises assemble intelligent components without the connective tissue required for autonomy. The

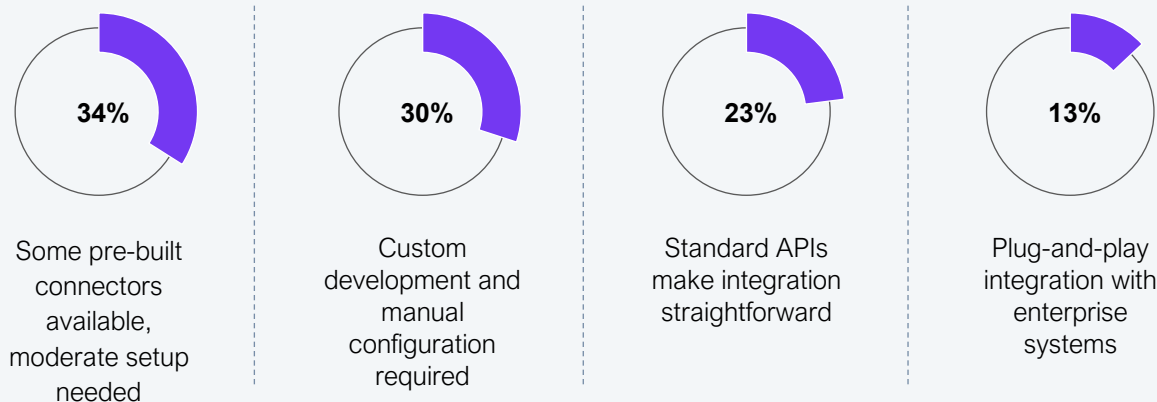
most advanced organizations are beginning to address this by building agent orchestration layers that govern communication, authentication, and logic across agents and systems.

## You cannot scale autonomy on brittle systems

The data converges on one conclusion: Enterprises are layering intelligence onto foundations still optimized for stability and control. Until leaders rebuild around real-time data, governed access, resilient security, and interoperable orchestration, they will scale complexity faster than capability.

### Exhibit 14: Integration complexity continues to limit agentic scalability

#### Which best describes your technology infrastructure for agentic AI?



Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

# The real barrier to agentic is not code; it's culture

As agentic systems gain autonomy, the limiting factor shifts decisively from technology to how humans and organizations are structured around it. The data does not point to mass resistance or fear of AI. It points to something more subtle and more consequential: Enterprises have not yet redesigned how authority, accountability, and trust work when machines begin to act.

Culture is not blocking progress. It is signaling unresolved design choices. Where boundaries are unclear, confidence erodes. Where accountability is ambiguous, autonomy stalls.

## Resistance reflects ambiguity, not rejection

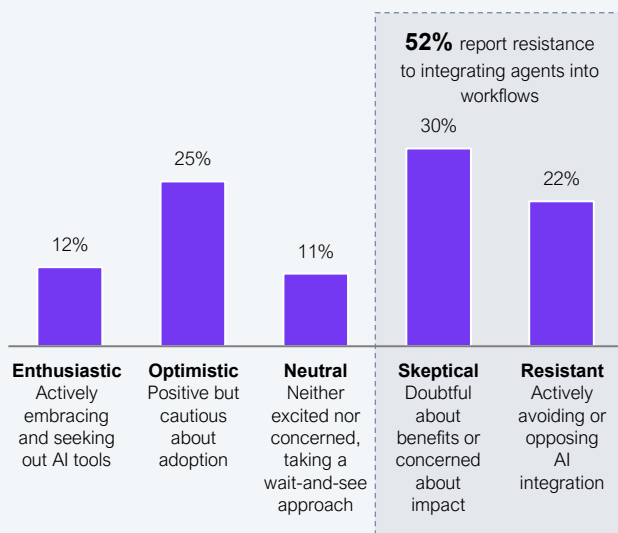
Workforce hesitation is real, but it is not ideological. Fifty-two percent (52%) report

resistance to integrating agents into workflows (see Exhibit 15), yet this resistance aligns closely with earlier signals around fragmented ownership, unclear governance, and limited delegation. Employees are not rejecting agents. They are questioning how responsibility is shared when outcomes are produced jointly by humans and systems.

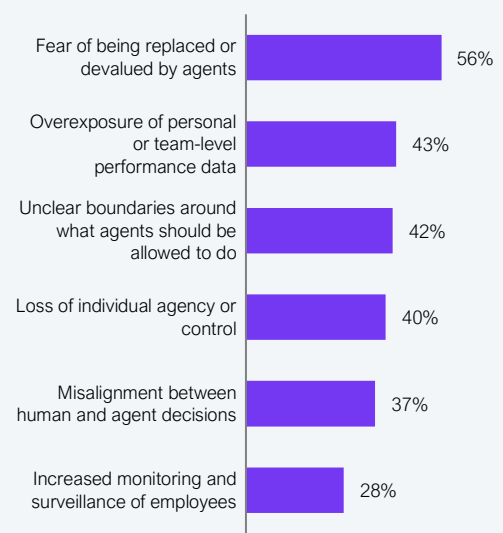
This distinction matters. Resistance is not a change-management problem to be solved with messaging. It is a design problem rooted in unclear escalation paths, poorly defined roles, and uncertainty about how decisions will be judged.

**Exhibit 15: More than half (52%) are resistant to agent integration in workflows**

### How are employees reacting to agent integration into workflows?



### What workforce concerns are most top-of-mind as visibility into work patterns increases with agentic AI?



Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

## Trust is conditional, and it is negotiated through visibility

Leaders express confidence that trust can be maintained as agentic AI scales. At the same time, 80% agree that agents need access to detailed employee performance and activity data to optimize workflows (see Exhibit 16). These two


signals coexist, but they create a cultural tension that enterprises have not fully resolved.

Autonomy requires visibility. Trust requires legitimacy. When employees do not understand how data is used, who sees it, or how it influences decisions, visibility feels like surveillance. Culture is shaped not by the presence of data, but by whether its use is transparent, bounded, and defensible.


### Exhibit 16: Workforce concerns center on privacy, control, and unclear boundaries

To what extent do you agree with the following statements about agentic AI?

**We can maintain employee trust while scaling autonomous AI systems.**



**For agentic AI to optimize workflows, it needs access to detailed employee performance and activity data.**



Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

## Leaders are cautious about delegation, and the organization feels it


Enterprises often frame hesitation as a workforce issue, but leadership caution plays an equal role. Executives remain reluctant to delegate complex or high-consequence decisions to agents (see Exhibit 17). This matters because autonomy

is learned through use. If leaders design systems that require constant human override, they reinforce containment even as adoption expands.


Autonomy cannot mature in an environment where authority is never fully granted. The result is a paradox: agents are deployed, but responsibility never truly shifts. Over time, this undermines confidence on both sides of the human-agent relationship.

### Exhibit 17: Leaders' own fears mirror those they project onto employees

#### Which parts of your work would you be willing to delegate to an autonomous agent today?



#### What holds you back from delegating more of your work to an agent?



Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

## Collaboration lacks rules, and culture absorbs the cost

Cultural readiness shows up in operational basics. Forty percent (40%) report having no formal training or guidelines for human-agent collaboration, and 37% say role boundaries are

unclear (see Exhibit 18). Without explicit handoffs, escalation rules, and accountability frameworks, collaboration becomes inefficient and fragile.

This is how trust erodes in practice, not through a dramatic failure, but through repeated moments of uncertainty about who owns the outcome when something goes wrong.

### Exhibit 18: The biggest pain points lie in ineffective human-agent collaboration training

#### What are the biggest challenges in human-agent collaboration in your organization?



Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

## Workforce strategy remains defensive, not redesign-led

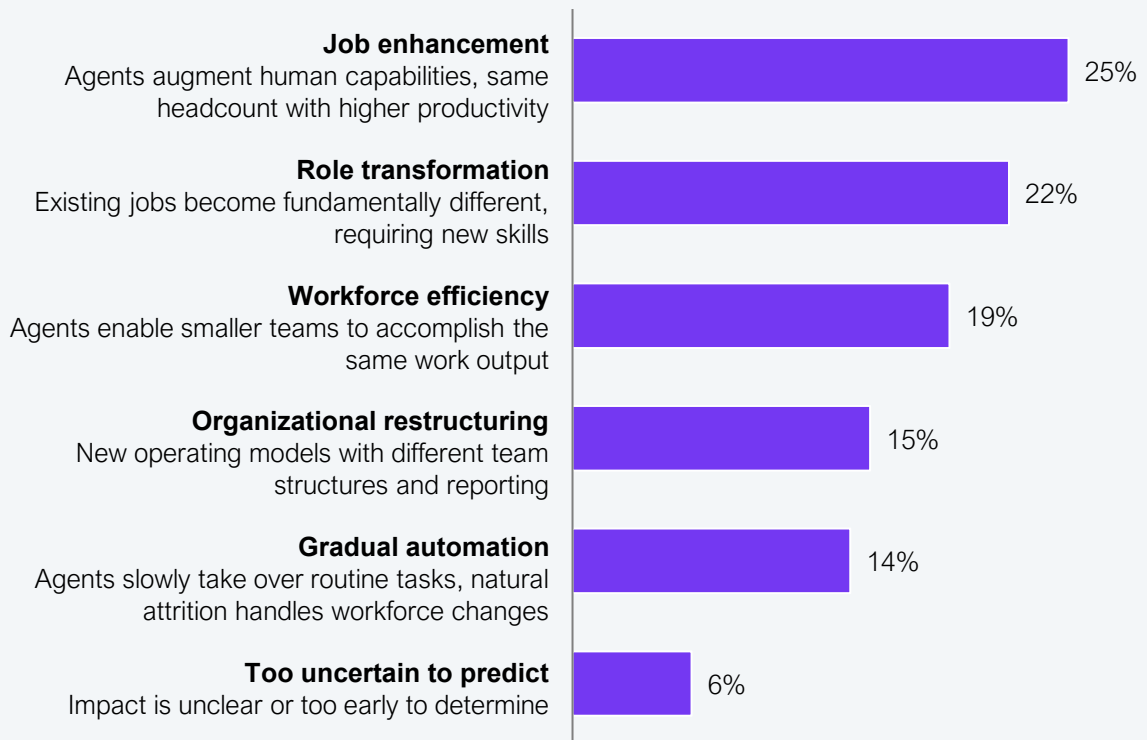
Most enterprises still describe their approach to agentic AI as job enhancement rather than workforce redesign (see Exhibit 19). This framing lowers anxiety, but it also constrains ambition. When the goal is to preserve existing roles,

organizations avoid rethinking how work should be decomposed, reassigned, and governed in an autonomous environment.

More mature organizations are beginning to treat agents as a way to shift humans toward judgment, synthesis, and exception handling. That shift requires intentional role design, not reassurance.

### Exhibit 19: Most leaders see agentic AI as job enhancement, not organizational reinvention

Which scenario best describes your organization's long-term vision for agentic AI and work over the next three to five years?



Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

## Operating models are beginning to bend around intelligence

Even with today's caution, leaders anticipate structural change. Nearly one in three expect to move toward process-centric operating models, and 21% envision network-based models where AI orchestrates talent and resources dynamically (see Exhibit 20). These expectations signal an emerging recognition that autonomy ultimately reshapes structure, not just tasks.

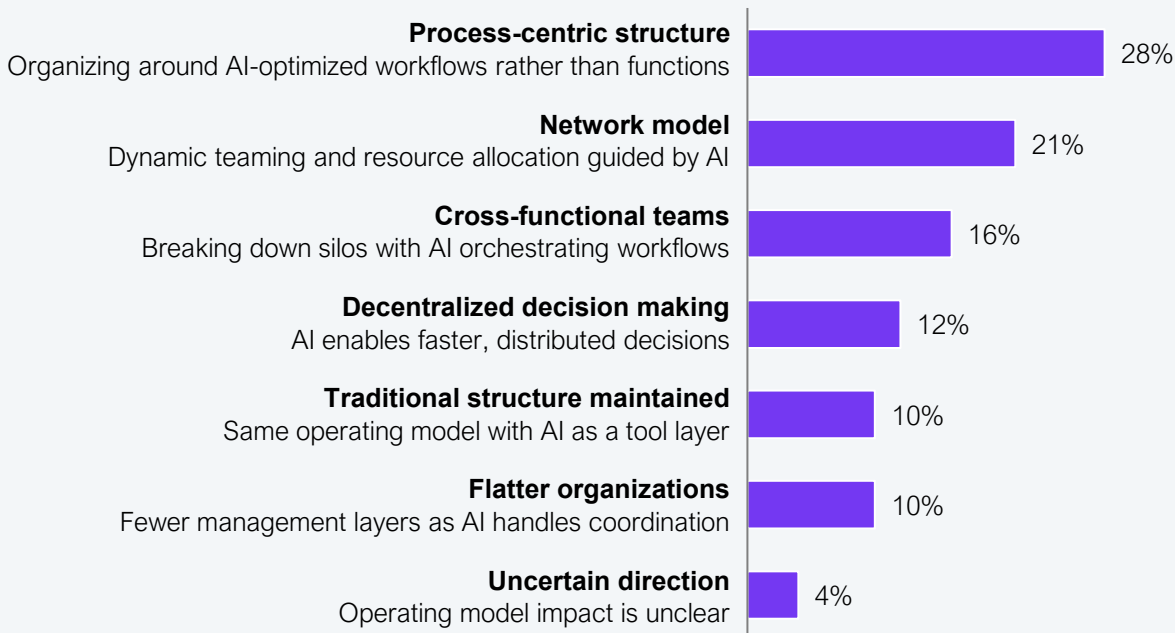
As intelligence becomes embedded in execution, the enterprise will be defined less by hierarchy and more by how work flows and decisions escalate.

## Culture is the pacing factor for autonomy

Technology is ready to move faster than organizations are comfortable with. The real constraint on agentic scale is not compute or model sophistication. It is whether enterprises can establish a credible human operating model for autonomy, one with clear boundaries, explicit accountability, and trust that is earned through design, not assumed.

### Exhibit 20: Enterprises anticipate a shift from function-based hierarchies to process- and network-based operating models

#### What operating model changes do you anticipate from scaling agentic AI?



Sample: 505 Global 2000 enterprise decision makers  
Source: HFS Research, 2026

## Autonomy will be built in workflows, not workarounds

The next leap in enterprise performance will not come from deploying more agents. It will come from redesigning how work moves. Intelligence cannot reach its potential inside operating models built for supervision, manual handoffs, and exception-heavy workarounds. Autonomy emerges when workflows themselves are reengineered to support coordinated action.

What separates incremental progress from real autonomy is not the sophistication of models, but whether enterprises are willing to redesign execution around intelligence rather than retrofit intelligence into legacy processes.


## Most enterprises are still adapting workflows to fit agents

The dominant deployment pattern mirrors earlier automation waves. Thirty-seven percent (37%) are layering agents onto existing workflows, and another 32% are making limited process optimizations to accommodate new capabilities (see Exhibit 21). Only one in five is redesigning workflows explicitly for autonomy.

This approach produces quick wins but shallow change. Agents improve speed and consistency inside old structures, but the underlying logic of execution remains unchanged. Work still flows through fragmented systems, unclear decision rights, and human-centric escalation paths.

### Exhibit 21: Most enterprises are still layering agents onto existing workflows rather than redesigning for autonomy

How would you describe your organization's current approach to process transformation with agentic AI?



	Exploring	Emerging	Scaling	Pioneering
We're redesigning or creating workflows to optimize the value of autonomous agents	22%	14%	29%	48%
We're optimizing existing processes to fit agentic capabilities	33%	29%	38%	28%
We're layering agents on top of existing workflows	35%	43%	28%	21%
We haven't made any structural changes to processes yet	10%	13%	6%	3%

Sample: 505 Global 2000 enterprise decision maker  
Source: HFS Research, 2026

## Redesign and orchestration are the real accelerators of scale

The data is clear on what actually moves the needle. Thirty-five percent (35%) cite redesigning core processes to support agent autonomy as the single most effective accelerator of adoption, followed by creating orchestration roles (30%) and embedding AI into broader transformation programs (26%) (see Exhibit 22). These are not technology upgrades. They are operating model decisions.

Enterprises that move fastest are treating workflows as living systems, instrumented, governed, and continuously improved. Orchestration becomes the connective tissue that allows agents and humans to coordinate across systems and functions without constant supervision.

## Process is becoming the system of record for autonomy

As autonomy expands, traditional systems of record lose primacy. What matters most is not where data is stored, but how decisions are sequenced, escalated, and audited across workflows. Process becomes the system of record that defines who can act, when intervention is required, and how accountability is enforced.

This is the quiet but profound shift underway. Enterprises cannot automate their way to autonomy. They must redesign execution itself.

### Exhibit 22: Process redesign and orchestration roles are the strongest accelerators of agentic AI adoption

#### What internal actions have helped your organization accelerate adoption of agentic AI?



Sample: 505 Global 2000 enterprise decision maker  
Source: HFS Research, 2026

## The operating model reset: How to join the 6%

Agentic AI is already moving from experimentation to real execution, but most enterprises are trying to scale it inside operating models built for control, approvals, and fragmented accountability. That is why progress stalls at a contained scope, even when capability improves.

This research points to a clear dividing line. The 6% that are scaling enterprise-wide scope are not just deploying more agents. They are redesigning how work gets done, who owns outcomes, how decisions escalate, and how leaders maintain confidence through visibility and evidence. The rest are scaling activity, not autonomy.

### What leaders should do next

#### **Pick one cross-functional workflow and rebuild it for delegated execution.**

Redesign handoffs, decision points, exceptions, and escalation so agents can act across the workflow, not just inside a task. If the scope cannot expand safely in one workflow, it will not expand anywhere.

#### **Assign outcome ownership for agent-run decisions, then enforce it.**

Name who owns the results when agents act, not just who owns the platform. Define intervention rights, escalation paths, and accountability when things go wrong, then bake it into governance.

#### **Build orchestration roles as the connective tissue.**

Process redesign and orchestration roles are the strongest accelerators (and that orchestration is what aligns agents, humans, and accountability across workflows).

#### **Make real-time data access and observability a gating condition.**

If agents cannot see the state of work in real time, they cannot own decisions. Prioritize data flow, monitoring, and audit trails before increasing agent authority.

#### **Redesign governance for systems that act.**

Move from static approvals to continuous controls, behavior monitoring, and post-decision auditability. Autonomy requires runtime governance, not governance theater.

#### **Define the rules of engagement for human-agent collaboration.**

Clarify when humans intervene, what exceptions look like, and how judgment is exercised. Train for oversight and exception handling, because that is where trust is built in practice.

#### **Stop counting agents and start measuring scope and decision quality.**

Track where agents operate, what decisions they own, how often they escalate, and how outcomes change. If you cannot measure decision performance, you are scaling activity, not autonomy.

#### **Pressure-test today's "safe constraints" as tomorrow's permanent limits.**

Oversight norms, permission policies, and containment patterns harden quickly. Design for the autonomy you intend to live with, not the one that feels comfortable this quarter.

## The 90-day test that proves you are joining the 6%

In 90 days, you should be able to point to one workflow that is redesigned for delegated execution, with named outcome ownership, defined escalation rules, runtime visibility (audit trails and monitoring), and metrics showing scope expansion and decision performance. This is the report's maturity test in practice, because scope is the constraint, and only a minority reports enterprise-wide deployment today.

### The 90-day test

If you cannot redesign and run one delegated workflow with clear ownership, defined escalation paths, runtime visibility, and outcome-based measures within 90 days, the constraint is not the agent technology. It is operating model inertia.

Joining the 6% is not about ambition. It is about making autonomy legible, accountable, and scalable through an operating model reset.

**The Bottom Line: If you scale agents before you redesign decision rights, data flow, and accountability, you will not get autonomy; you will get faster fragmentation.**

The enterprises that get this right will not just deploy smarter systems; they will redesign how intelligence actually moves through the organization. Agentic maturity is defined less by the number of agents deployed and more by how clearly authority, escalation, and accountability are structured. Enterprises that institutionalize that design principle will convert experimentation into enterprise-wide execution change.

## HFS Research authors



**Dana Daher**  
Executive Research Leader

Dana Daher is an Executive Research Leader at HFS Research, spearheading research initiatives in emerging technologies and employee experience. With a unique blend of expertise in anthropology and IT, Dana leads cutting-edge research that shapes industry landscapes across various domains, including employee experience, Agentic AI, generative AI, diversity, equity, and inclusion (DEI), and sustainability.



**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.




**Hridika Biswas**  
Associate Practice Leader

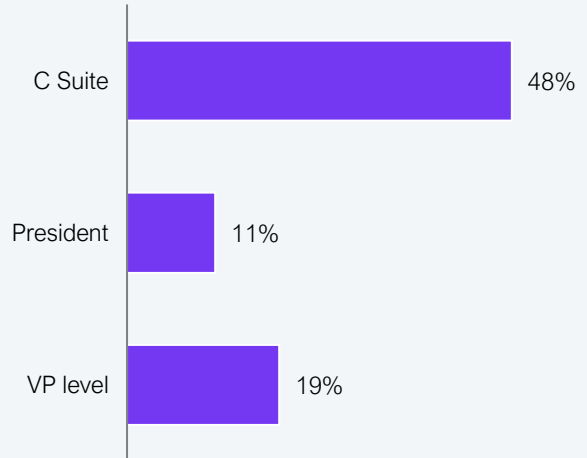
Hridika Biswas is an Associate Practice Leader at HFS Research. She joined the team in 2021 and her key coverage areas include business services (F&A, S2P), process intelligence, intelligent document processing, and automation. Her interests include analyzing how emerging tech enables enterprises to reach their transformation goals.

## Participant demographics (1/2)


### By Region



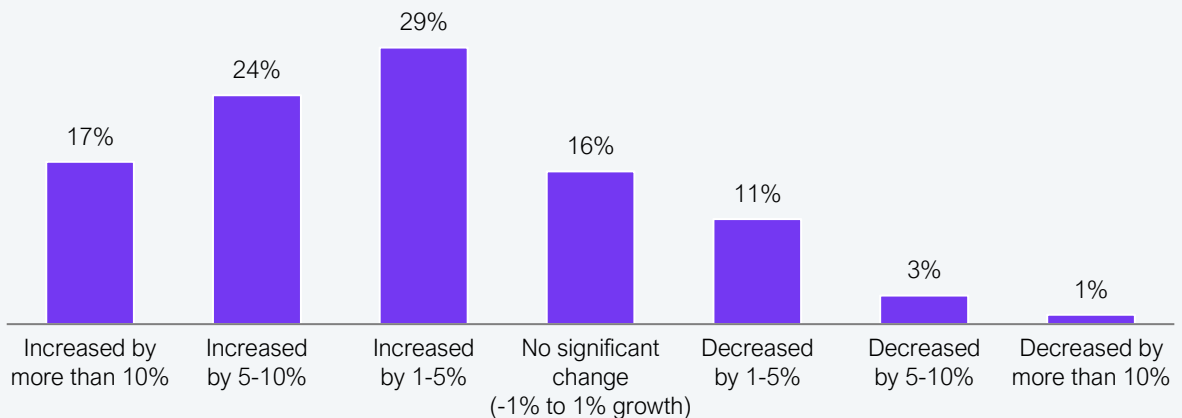
### By Role



### What is your organization's annual revenue range?




### How did your organization's annual revenue change compared to the previous fiscal year?




Sample: 505 Global 2000 enterprise decision maker  
Source: HFS Research, 2026

# Participant demographics (2/2)

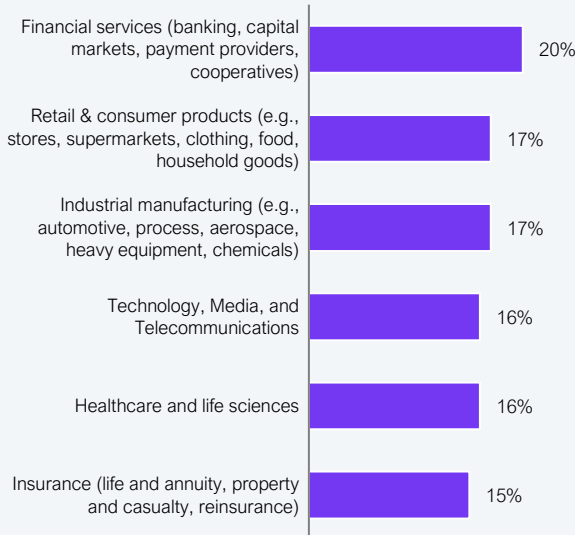
## What is your organization's current stage of Agentic AI adoption?



## Which of the following best describes your organization's current stage of GenAI adoption?



### By Industry



### By Primary Function



Sample: 505 Global 2000 enterprise decision maker  
 Source: HFS Research, 2026



## About Infosys

Infosys is a global leader in next-generation digital services and consulting. Over 320,000 of our people work to amplify human potential and create the next opportunity for people, businesses, and communities. We enable clients in 59 countries to navigate their digital transformation. With over four decades of experience in managing the systems and workings of global enterprises, we expertly steer clients, as they navigate their digital transformation powered by cloud and AI. We enable them with an AI-first core, empower the business with agile digital at scale and drive continuous improvement with always-on learning through the transfer of digital skills, expertise, and ideas from our innovation ecosystem. We are deeply committed to being a well-governed, environmentally sustainable organization where diverse talent thrives in an inclusive workplace.

Visit [www.infosys.com](http://www.infosys.com) to see how Infosys (NSE, BSE, NYSE: INFY) can help your enterprise navigate your next.



## About Infosys Topaz

Infosys Topaz is an AI-first set of services, solutions and platforms using generative AI technologies. It helps amplify the potential of humans, enterprises and communities to create value from unprecedented innovations, pervasive efficiencies and connected ecosystems.

## 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](http://www.hfsresearch.com)



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



[www.horsesforsources.com](http://www.horsesforsources.com)



[www.horsesmouthpodcast.com](http://www.horsesmouthpodcast.com)