



the results company

The State of Enterprise Agentic AI

Every piece in place. Every outcome by design.

A diagnostic briefing on where the ROI gap is widening, what's missing for most organizations, and what it takes to realize returns faster.



Table of Contents

EXECUTIVE SUMMARY	01
THE ROI GAP ISN'T A TECHNOLOGY PROBLEM	02
WHY MOST ORGANIZATIONS ARE STUCK	04
THE PROOF: WHAT THE RIGHT DESIGN DELIVERS	05
THREE MISTAKES THAT DERAIL AGENTIC AI	06
THE STANDARD: THREE NON-NEGOTIABLES	08
THE ART OF GETTING EVERY PIECE IN PLACE	09
QUESTIONS TO PRESSURE TEST YOUR INVESTMENT	11
FIND YOUR FASTEST PATH FORWARD	12

Executive Summary

Agentic AI has moved from concept to deployment faster than any recent enterprise technology. The investment has followed. For most organizations, the returns have not.

The gap between what agentic AI promises and what it delivers is now the defining problem for the leaders responsible for both. Many of the organizations generating significant and timely returns from agentic AI are using the same general capabilities available to everyone else. What separates them from the rest is deeper: the data foundation it operates against, the workflows it is embedded in, and the governance that holds it accountable.

This report examines those distinctions. Drawing on research from Deloitte, MIT, Gartner, Forrester, McKinsey, and Microsoft, alongside outcomes delivered in production environments, it lays out where the ROI gap is widening, the missteps putting most deployments at risk, and the non-negotiables shared by effective deployments.

The data and outcomes that follow are curated for leaders carrying the weight of the agentic AI mandate—the executives tasked with deploying agentic AI at scale and defending the investment to the board, without exposing the organization to risk it cannot afford.

The leaders poised to deliver meaningful returns are not waiting for the picture to come into focus. They are designing it.



THE ROI Gap Isn't a Technology Problem

The AI ROI gap is real and already measured.

85%

of organizations increased AI investment last year.¹



6%

Only 6% of those organizations saw any return.¹



95%

of GenAI deployments delivered zero P&L impact.⁹



Agents amplify.



Given ungoverned data, they produce outputs nobody trusts.



Dropped into workflows built for humans, they create friction instead of removing it.



Deployed without guardrails, they expose the organization to significant risk.

That gap between what the technology can do and what the environment can support separates the organizations generating measurable returns from the ones still running pilots.

Organizations making headway are already compounding the advantage.

The gap between investment and return is not inevitable. It is a function of how the deployment was designed.



**120% ROI &
\$24.2M**

in net present value from Microsoft Agentic AI Solutions.³

A January 2026 Forrester Total Economic Impact study, commissioned by Microsoft, found that enterprise deployments of Microsoft's agentic AI solutions delivered a 120% ROI and \$24.2M in net present value over three years, with payback in 15 months.



The Real ROI Gap

Waiting looks safe. It Isn't.

The instinct to let the technology mature or to run one more proof of concept before committing is understandable. It is also now one of the most expensive decisions an enterprise leader can make. Not because agentic AI is a trend that will pass. Because the organizations moving today are building structural advantages that expand year-over-year.

Microsoft and IDC's 2025 research defines Frontier Firms as organizations deploying AI across seven or more business functions. **Compared to slow adopters, Frontier Firms report better outcomes at a rate that is four times higher across:**



(Microsoft/IDC, Becoming a Frontier Firm, December 2025)

The question is not whether to adopt agentic AI. It is whether to adopt it in a controlled, structured way, or be forced into it later at greater cost and with fewer options.

The fastest path to payoff is the most deliberate one. Design is the differentiator.

Why Most Organizations are Stuck in Purgatory

Most enterprises pursuing agentic AI are not succeeding and not visibly failing. They are stuck in a middle state: a purgatory of ambitious deployments that launched, demonstrated technical capability, and then stalled before generating measurable returns.

95%

of enterprise AI deployments show no measurable P&L impact. Agentic AI is arriving into the same pattern.

(MIT NANDA, State of AI in Business 2025)



Deloitte's 2025 research captured the pattern directly.

60% of enterprises have now given most of their workforce access to sanctioned AI tools, up from under 40% the year before.²

25% have moved 40% or more of their AI experiments into production.²

(Deloitte, State of AI in the Enterprise, 2025)

Agentic AI intensifies this holding pattern. Recent MIT Technology Review Insights research found that



98% of senior technology executives believe agentic AI is required to escape “pilot purgatory,” and that traditional AI pilots fail because they stop at recommendation rather than execution.¹⁰

(MIT Technology Review Insights, sponsored by SoftServe, 2026)



A new failure mode has emerged alongside the old ones: agentic pilots that launch with full autonomy but without the integration, governance, and KPI alignment to support it.



Agentic AI does not eliminate purgatory. It exposes the foundational gaps that put organizations there in the first place, just faster, and with greater visibility, than any prior technology wave.



For organizations already stuck, the path forward is not more pilots. It is a different starting point.

The Proof: What the Right Design Delivers

These are outcomes already delivered for organizations and designed by HSO. To see more, visit hso.com/agentical

Every outcome below was delivered for a real organization, in a real production environment, on a foundation that was prepared before agents were deployed. Each is the result of a deliberate sequence, executed by a partner who held the full picture from day one and measured success the way the board does.

The pattern across every outcome is the same. **Foundation first. A partner designing for outcomes.** With that approach, the path from idea to impact becomes shorter, and more certain.



**98%
less**

Manual
Processing

Hospitality

A global organization managing operations across more than 600 locations faced a data migration and operational reporting cycle that consumed 160 hours of processing time per run. After HSO's deployment of a data migration agent on Microsoft Dynamics 365, that cycle runs in four hours.



**15,000
hours**

Back Every
Year

Retail Distribution

A luxury consumer goods brand scaling global operations needed AI-powered customer engagement that worked within strict data governance requirements. HSO's deployment returned 15,000 hours annually to the organization by automating fulfillment. Even better, they have now achieved 4.9 out of 5 customer satisfaction.



**Weeks
to
Minutes**

Professional Services

A professional services firm working in more than 30 countries faced a scaling problem. Generating client-specific reports was a manual, time-intensive process that could not keep pace with demand. HSO implemented a suite of AI-driven knowledge



**30%
less**
Manual
Effort

Financial Services

A financial services organization operating under significant regulatory scrutiny needed to address manual overhead in credit approval and fraud detection. HSO's deployment reduced manual effort by 30%, with security and compliance built into the architecture from the start. The wins began in week one, with 40 forged paystips detected.

Three Mistakes That Derail Agentic AI

By the time most organizations recognize these missteps, they've already paid for them.

Most agentic AI initiatives don't fail loudly. The pilot runs. The demo impresses. The steering committee approves the next phase. Then, somewhere between proof of concept and production, the momentum disappears. The outputs are not trusted. The workflows do not change. Six months later, leaders at the organization are back in the same conversation: more cautious, more skeptical, and further from meaningful returns than when they started.



Building Agents Before Foundations

Organizations identify a promising scenario, build or buy an agent, and then discover in production that the data the agent needs is not structured, not trusted, or not accessible in the way the agent requires. Or they find that the agent works in an ineffective step of a workflow. MIT NANDA found that only 5% of enterprises have AI tools integrated into workflows at scale, and seven of nine major industry sectors show no real structural change from AI adoption despite widespread investment. The activity and interest continues to grow, but the mismatch between investment and returns remains.

Treating agents as add-ons

An agent that sits alongside a business process is optional. An agent embedded inside a redesigned process is how work actually gets done.

This is where most deployments fall short. Organizations layer agents on top of existing workflows, built on inconsistent data, and expect them to compensate for both. Agents work with what they are given. Ungoverned data produces outputs nobody trusts. Workflows that require a human to navigate four systems and three approvals to act on an agent's recommendation deliver nothing of value.

High-performing organizations take a different approach. They redesign workflows first—then embed agents into them. Most organizations are doing the opposite, calling it AI adoption. The result is predictable: tools that don't learn, don't retain context, and don't integrate with how work actually happens.

The agents delivering returns are embedded in redesigned workflows—moving businesses from human-by-default, where every action must be initiated and approved by a person, to human-by-exception, where only the decisions that matter reach one.



Governance added after the fact

Leaders who approach governance as a Phase 2 concern pay twice: once to fix the problem, once in the credibility spent in the meantime.

The risks at the governance layer are not always obvious: vulnerabilities at the integration layer that only surface at scale, audit trail gaps that become material when a regulator asks a question the organization cannot answer, policy conflicts the agent resolves incorrectly because nobody defined the governing logic before deployment. Gartner predicts more than 40% of agentic AI projects will be canceled by the end of 2027, driven by escalating costs, unclear business value, and inadequate risk controls. Governance failures show up in all three.

The Standard: Three Non-Negotiables

Observable, measurable, and already proven across the organizations delivering real returns.

The research describes the problem from multiple angles: the ROI gap, the pilot-to-production drop-off, the workflow integration failure. What it doesn't do is prescribe what it takes to succeed. That pattern has to be observed in the field, across enough deployments, in enough conditions, to separate what actually works from what only sounds like it should.

The variables that determine whether agentic AI delivers returns are neither surprising nor complex. The platform it runs on. The data it operates against. The governance that holds it accountable. Each is a decision. The three together provide the architecture.



Three Pillars for Scalable Agentic AI



01

A platform you can trust

Agentic AI that runs on infrastructure the organization does not fully control is agentic AI the organization cannot fully defend.

In practice, that means deployment inside the organization's own cloud environment, grounded in approved and verified knowledge sources, governed by existing security and compliance frameworks.

Data does not leave the environment.

Outputs can be traced, audited, and explained to anyone who asks.



02

Data that is agent-ready, not just gathered

Most enterprise organizations have data. What they often lack is governed data: structured, trusted, consistent across systems, and accessible in the way agentic AI requires. The distinction matters because agents magnify whatever they are given.

Clean, governed data produces outputs that can be trusted and acted on.

Ungoverned data produces outputs that introduce risk faster than they create value.

Gartner predicts organizations will abandon **60% of AI projects in 2026 due to insufficient data quality.**

The MIT research confirms how early this shows up: **only 5% of enterprises have AI tools integrated into workflows at scale, with data readiness a primary cause.**



03

Security & Compliance by design.

The organizations reaping returns from agentic AI did not add security and compliance onto their agentic AI deployments. They built their deployments around them.

That framework enables speed to results, because it builds successes and eliminates the need for retrofits in production.

Deploying an agent without designed-in governance is the equivalent of hiring someone and ignoring the onboarding, the defined role, planned performance reviews, and KPIs. The agent will act based on guardrails and expectations. Skip those and the cost of unravelling and reworking later is significantly higher than the cost of designing it correctly at the start.

The Art of

Getting Every Piece in Place

The best agentic AI deployments are design disciplines.

The difference between deployments that are delivering returns and those that are not is a function of how each was approached and by whom it was designed. The organizations gaining ground treat agentic AI as a design problem. They begin with the outcome and work backwards: to the foundation, the governance framework, the workflow, and the specific agents required to deliver it. They sequence correctly, starting with strategy before solution, governance before deployment, outcomes before activity.

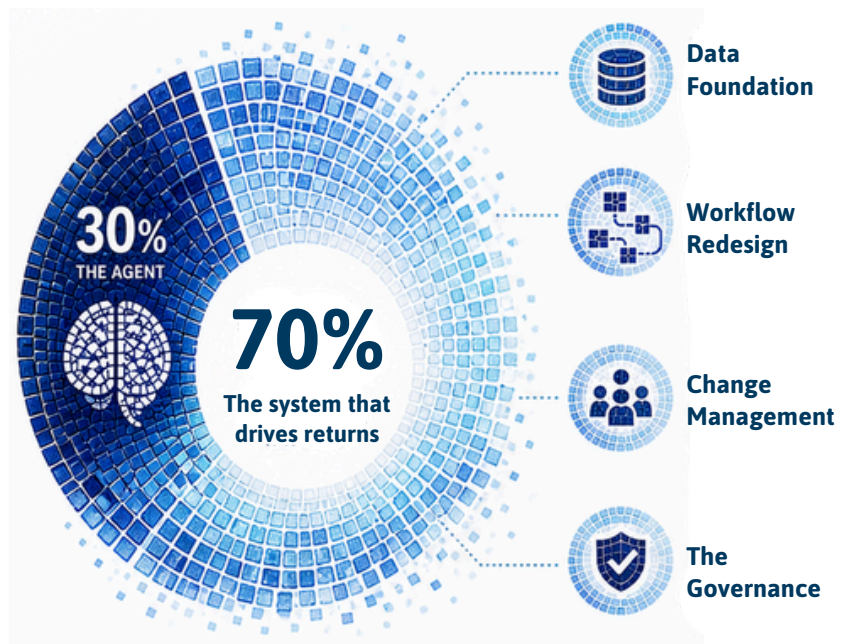
“The agent is
30%
of the solution.”

The data foundation, the workflow redesign, the change management, the governance — that’s the other 70%.

Most organizations are skipping it because it’s harder to see and harder to sell internally than the agent itself.”

— Touseef Zafar
HSO Global Service Line Technology Lead, AP

Forrester's 2026 TEI study of Microsoft agentic AI deployments treats foundational elements—governance, security and compliance, platform tools and data buildout—as a distinct, ongoing investment category alongside agent development itself.³



The 30% gets approved.

The 70% determines if it works.

What more than three decades of Microsoft Implementation experience delivers.

HSO has been navigating every major wave of Microsoft innovation for more than 30 years, as a Microsoft-exclusive partner by deliberate choice. That focus means something specific in practice. HSO's consultants know not only what Microsoft's platforms can do in a controlled demonstration, but also what they do under enterprise conditions—at scale, in complex regulated environments, and across multiple geographies and business units simultaneously.

It also means HSO has seen every version of the problems the current moment presents. Including:



Data that isn't ready



Poorly connected workflows



Governance that arrives too late



Partners who disappear after the pilot

HSO has built a methodology specifically designed to interrupt these patterns before they become expensive.

HSO holds all six Microsoft AI Cloud Solution Designations and has been recognized by Forrester as a Strong Performer in the 2026 *Forrester Wave™: Microsoft Business Applications Services, Q1 2026*.

“HSO’s approach emphasizes modern, scalable transformation, evolving from Copilots to autonomous agents via its AI Factory and agentic AI roadmap and co-engineering with Microsoft (e.g., PayFlow Agent & Supplier Relationship Management).”

Biswajeet Mahapatra | Principal Analyst, Forrester, as quoted in the 2026 Forrester Wave™: Microsoft Business Applications Services, Q1 2026

The result is a path from idea to impact that is faster than expected, because the full picture was addressed up front. Agents that go live are the agents that were designed to deliver, in an environment designed to support them: inside Microsoft 365, Dynamics 365, and Teams, where work already happens.

Questions to Pressure Test Your Next Investment

These three questions help surface the most costly pitfalls in agentic AI deployment.

The organizations that get agentic AI right do not always start with more resources or better technology. They start with a sharper strategy. These questions apply whether you are still defining your strategy or already in production.

01

Can you define what success looks like in specific, measurable terms?

Not "we want to reduce manual effort." A specific outcome, with a number attached, and a realistic timeframe. This is the definition that drives every decision that follows: which agents get built, which data needs to be in order, which workflows need to change, and how you will know whether any of it worked.

02

Have you examined the state of the data your agents will operate with?

Data that works for reporting does not automatically work for autonomous agent action. The divide between "good enough for analytics" and "ready for agents" is where a significant number of deployments run into trouble. An honest assessment of data readiness before deployment is the strongest predictor of whether your agents will produce outputs people trust and adopt. Agentic AI surfaces data weaknesses faster than any system that came before it. The organizations that treat that exposure as an asset come out ahead.

03

Is governance part of your deployment architecture?

If you have not deployed yet, governance designed from the start makes deployment faster, not slower. It removes the need to retrofit security and compliance into an environment that was not built to accommodate them. If you are already in production, the question is what it will cost to address governance gaps now versus what it costs when a regulator, an auditor, or an incident forces the timing.

Find Your Fastest Path Forward

The distance from where you are to results that matter is shorter than most organizations expect.



For one municipal government, the HSO-designed pathway meant **going live in eight weeks and resolving 95% of citizen inquiries without human intervention.**



For a 10,000-person enterprise modeled by Forrester, it meant **\$24.2M in net present value and 15-month payback on investment.**

Same principle. Foundation first, agents second.

The path forward looks different depending on where you are today.



Ready to move now

Deploy immediately on a governed, enterprise-ready platform.



Looking for fast proof

Launch proven agents that deliver measurable returns in weeks.



Need to clarify first

Start with strategy to ensure everything that follows is faster and more certain.



HSO's Agentic AI Advisory Service

Take the guesswork out of your next move.

It includes the full picture: a structured engagement that assesses your environment, identifies where agents will provide the most impact, and designs the governance and data foundation to support them—so every agent that follows is built to deliver.

 **design your results**

References

- ¹ Deloitte. *AI ROI: The Paradox of Rising Investment and Elusive Returns*. October 2025. Survey of 1,854 executives, Europe and Middle East.
- ² Deloitte. *State of AI in the Enterprise*. 2025. Based on a survey of more than 3,200 business and IT leaders with direct involvement in their companies' AI initiatives.
- ³ Forrester Consulting. *The Total Economic Impact™ of Microsoft's Agentic AI Solutions*. Commissioned by Microsoft. January 2026. Based on interviews with 8 decision-makers across 6 organizations and a survey of 420 respondents with experience using Microsoft's agentic AI solutions.
- ⁴ Forrester. *The Forrester Wave™: Microsoft Business Applications Services, Q1 2026*.
- ⁵ Gartner. *Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027*. 2025.
- ⁶ Gartner. *Lack of AI-Ready Data Puts AI Projects at Risk*. February 2025.
- ⁷ McKinsey & Company. *The State of AI*. 2024–2025.
- ⁸ Microsoft and IDC. *Becoming a Frontier Firm: Unlocking the Business Value of AI*. December 2025.
- ⁹ MIT NANDA. *State of AI in Business 2025: The GenAI Divide*. July 2025. Based on 52 structured interviews, 153 senior leader surveys, and analysis of 300+ public AI initiatives.
- ¹⁰ MIT Technology Review Insights. *Redefining the Future of Software Engineering*. Sponsored by SoftServe. 2026. Based on a survey of 300 CIOs, CTOs, CDAOs, Chief AI Officers, and other senior technology executives across seven industries and six countries, fielded December 2025–January 2026.



Next steps



Learn More About The Art of Agentic AI



HSO's Agentic AI Advisory Service



Agentic AI in Action - Customer Success stories

