



State of Enterprise AI Adoption

September 2025




Table of Contents

3 Executive Summary

5 Methodology

6 The Evolving Value of Applied AI

11 State of the Market

14 Use Cases Are Moving Closer to Revenue

15 IT: Productivity Gains vs. Operational Depth

16 Back and Middle Office: Balancing Risk and Efficiency
in Core Operations

17 Front Office: Experimentation at the Customer Edge

18 Industry-Specific Use Cases: Variation Emerges

19 The Enterprise AI Ecosystem

21 AI Platforms

22 Enterprise Software

23 Service Providers

24 AI Pricing: Models Meet Markets

28 Requirements for Achieving Value with AI

29 Data by Design

32 Adoption Accelerators

33 AI Governance

34 Key Contributors

34 Related Research




Executive Summary

Executive Summary

In our third annual study of the AI market, we've broadened from primarily studying the impact of generative AI to now include agentic AI and the overall AI market. AI hype is still rampant, but our focus is only on enterprise adoption and value. To understand the state of adoption today, we asked an audience of **400 senior AI decision makers to profile their top 3 most funded use cases**. Our research found that 1 in 3 of the most funded uses cases are in production today, but only about one in four initiatives meets revenue impact expectations, and broad cost savings are elusive.

We had previously observed that the timelines for adoption had been slipping in the face of complexity, data integration challenges, measurement issues and tooling maturity. Now, as the number of production use cases climb, we see that the large variety of unique use cases is also climbing, and this provides a path for other companies to be fast followers. The relatively low convergence in the use case data shows just how much opportunity there is.

In this report, we'll dive into what enterprises are currently trying, where they are currently getting value, and what our findings suggest about the best practices needed to increase adoption, scale and return on investment.



1

The number of use cases in production has doubled since 2024

This doubling shows us that AI adoption is accelerating, underscoring faster progress in scaling. At the same time, the use case portfolio is diversifying. Enterprises are pursuing a broader, more varied set of initiatives with a long tail of unique, business-specific use cases. This mix reflects a dual reality: scale is increasing for proven functions, while experimentation continues at the edge.

2

Business outcomes lag AI ambition

Enterprises are scaling AI faster, but results remain uneven. The most consistent benefits are in compliance, risk management and quality control, where AI strengthens established processes. By contrast, growth and cost-reduction outcomes are underdelivering: only about one in four initiatives meets expectations for revenue impact, and broad cost savings remain elusive.

3

AI use cases are moving closer to revenue-generating activities

Enterprises are shifting their focus from efficiency-oriented experiments to functions that directly support growth. In 2025, leading use cases are centered on customer relationship management (CRM) automation, sales enablement, forecasting and lead capture — a change from 2024. While these initiatives don't yet create new revenue streams, they are embedded in processes tied to pipeline growth and customer acquisition, signaling that AI is being steered into higher value functions.

4 Industry priorities begin to diverge

Horizontal use cases (chatbots, copilots, and forecasting) dominate across industries, but vertical priorities are surfacing. Healthcare emphasizes clinical note generation, manufacturing and retail focus on forecasting, and banking targets lead capture and content generation. These sector-specific use cases are fewer but differentiating, requiring specialized data and domain expertise.

5 Autonomy-level pricing will emerge

Two pricing models dominate in 2025: software-style deals, where enterprises pay licenses and tokens but absorb hidden costs for training, monitoring and infrastructure; and transaction-based contracts, where agents perform work once priced per claim or ticket. ISG proposes new autonomy-level pricing (ALP) to bridge the gap. ALP tags resource units with autonomy levels so charges reflect how work is done, not just what is delivered.

6 AI value requires strong foundations


Ambition alone doesn't scale AI. Sustainable impact depends on three foundations: data by design – trusted, contextual, real-time information; adoption accelerators – workflows and skills that embed AI into decision-making; and governance – adaptive controls that manage risk without delays. Where these elements align, pilots move to production. Where they don't, value stalls.

METHODOLOGY

This ISG State of Applied AI report outlines and analyzes the most recent developments in applied AI through the lens of enterprise use cases. We studied 1,200 use cases derived from surveys of enterprise AI leaders and practitioners in large global organizations of 1,000 employees or more, including:

- A wide range of industry sectors including banking, insurance, manufacturing, IT, healthcare, retail, energy, media and public sector.
- Front and back office operations, IT and industry-specific use cases. Over 300 types of industry- or function-specific use case were reported and aggregated.
- A cross-section of industry functions including sales, marketing, product management, operations, finance, procurement, customer service and IT. Over a half of respondents were in C-level roles.
- The use cases studied represent a total of \$2.6 billion in AI spending, impacting 3 million employees.

ISG influences over \$200 billion of technology spending annually, which provides us with unprecedented insights into the intersection of large enterprises and the AI ecosystem. We use this knowledge, combined with the data in this and previous studies, to provide insights to help enterprises make informed decisions and improve the impact of their AI initiatives.

The background of the slide features a large, abstract, flowing shape in shades of dark teal and blue, resembling a liquid or smoke effect. Below this shape is a horizontal bar composed of several colored segments: dark blue, light blue, grey, lime green, dark blue, dark blue, light blue, grey, and lime green. The title 'The Evolving Value of Applied AI' is written in a large, white, serif font across the bottom half of the slide.

The Evolving Value of Applied AI

The Evolving Value of Applied AI

Artificial intelligence (AI) in the enterprise hasn't been a single breakthrough. It has unfolded in eras that layer on each other, expanding what the technology can do and how it creates value.

You can read this evolution in two ways. By technology: predictive, generative, agentic. Or by value: prediction, creation, orchestration. Here, we use both lenses with a focus on enterprise value.



Era of Prediction: Traditional AI

Built on structured data and careful models, this era was about recognizing patterns and making reliable calls.

When the data was clean and curated, results were precise. The value was confidence and scale, mostly inside narrow, well-defined functions.

Many organizations still begin here and then layer in generative and agent capabilities as trust, data and governance catch up.

Predictive AI examples: Fraud detection in banking; predictive maintenance in manufacturing; product recommendations in retail.



Era of Creation: Generative AI (GenAI)

Then came creation. AI could draft text, summarize calls, write code and synthesize insights from piles of unstructured information.

Teams moved faster and covered more ground, even as new risks showed up – hallucinations, biases and outputs that didn't always tie neatly back to source data. Value now included speed, breadth and imagination – as long as governance could keep pace.

GenAI Examples: proposal acceleration for sales; call center summarization; code-generation assistants.



Era of Orchestration: Agentic AI

Now AI is starting to act. Agentic systems connect earlier capabilities with agents that operate across tools and data, make decisions and carry them out.

The shift moves from single functions to a system-of-systems. Like microservices, agents create value at the intersections – where supply meets demand, finance meets operations and customer experience meets back-office execution.

Agentic AI Examples: supply chain agents adjusting purchase orders in real time; finance agents reconciling invoices; IT agents resolving incidents before escalation; customer-facing agents coordinating offers across CRM and marketing.


Unlocking the Value of AI

Technology evolves in eras, but value only shows up when enterprises are ready. Four essentials matter to achieving value:

- **Data that fits the task.** Data that is reliable in type, volume and context, structured for prediction, multimodal data for creation, and real-time and governable for orchestration.
- **Governance that adapts.** Oversight that enables speed and confidence, not just compliance checklists.
- **Enterprises built to collaborate with AI.** Workflows and decision rights that treat AI as a partner, not a bolt-on.
- **A workforce that's prepared and enabled.** People equipped to guide, monitor and extend what AI can do.


Some organizations still get the most from traditional, predictive AI. Others are unlocking creativity with generative tools. A smaller set is beginning to orchestrate with agentic systems. The roadmap isn't only about technology's evolution – it's about readiness to apply it.

So where are enterprises **today** on this curve?



Requirements, Risks and Rewards Vary by AI Tech


Traditional AI



Structured Data
↓
Clear Answers

- AI Behavior:** Rules-based, predictable
- Data Type:** Curated, structured datasets
- Use Cases:** Fraud detection, reporting, forecasting
- Governance:** Pre-cleaned data, manual validation, clear data ownership
- Outcome:**
- Trusted insights
 - Narrow, reliable use
 - **Low risk**

Generative AI



Big Data
↓
New Content, New Risks

- AI Behavior:** Creates new content (text, images, etc.)
- Data Type:** Massive, mixed
- Use Cases:** Chatbots, summarization, ideation
- Governance:** Less oversight, risks of hallucinations & bias, post-output reviews
- Outcome:**
- Increased potential
 - Decreased control
 - **Medium-high risk**

Agentic AI



Big Data
↓
New Content, New Risks

- AI Behavior:** Makes real decisions and executes tasks
- Data Type:** Dynamic, raw, real-time application data
- Use Cases:** Supply chain adjustments, autonomous systems
- Governance:** Often bypasses pipelines, no time for manual validation, needs live, in-line validation
- Outcome:**
- Business-impacting decisions
 - Invisible data dependencies
 - **High risk of operational failure**

How Enterprises Are Using AI Today

While looking at the eras shows how value expands, the day-to-day picture is grounded in what teams can reliably do right now. Most initiatives still center on AI's **predictive** strengths – **classification and recognition (33% of use cases)** and **optimization (26%)** – because these are the most mature capabilities and depend on structured data and clearly-defined outputs.

Generation and summarization (17%) are growing categories, often in supporting roles like drafting sales content or summarizing service interactions. Agent-like categories – **augmentation, automation and interaction were less than 10% each** – are emerging but remain a smaller share today.

Enterprises are starting from proven predictive capabilities and layering in generative and agentic enhancements as data readiness and governance improve.


This mix also helps explain outcomes: the strongest production results cluster around **compliance and risk**, while **revenue growth and cost reduction** remain harder to realize at scale. The market's center of gravity hasn't shifted as far as the hype suggests – yet.

Looking Ahead

Each era layers on top of the last. Prediction remains vital within creation. Creation fuels orchestration. And orchestration will keep evolving, reshaping how enterprises design systems, manage data and unlock new forms of value.

In the pages ahead, we look at where value is repeatable today, what's at the frontier, and what isn't working yet when AI is applied inside the enterprise.

Classifying the AI Use Cases



Source: ISG Market Lens, State of AI, 2025; n=1,200 use cases



Executive Summary

The Evolving Value of Applied AI

State of the Market

The Enterprise AI Ecosystem

Models Meet Markets

Requirements

Key Contributors

State of the Market

State of the Market

In 2025, 31% of the use cases studied reached full production, which is double the amount compared to our 2024 study. Additionally, we had previously observed that the timelines for adoption had been slipping in the face of complexity, data integration challenges, measurement issues and tooling maturity. Now, as the number of production use cases climb, we see that the long-tail of unique use cases is also climbing, and this provides a path for other companies to be fast followers for a growing number of use cases.

On the flip side, organizations entered 2025 with expectations that AI would cut costs and boost productivity, but as can be seen on the right, these outcomes are underdelivering. Only one in four initiatives is achieving expected ROI on growth, while just half are delivering on efficiency gains. Where AI is outperforming is in less flashy but critical areas like compliance, risk and quality, where it amplifies established processes.


\$1.3M

spent per use case on average, to date

31%

of prioritized use cases are in production

AI Use Case Metrics: Performance Relative to Expectations



Source: ISG Market Lens, State of AI, 2025; n=1,200 use cases

Use Cases Are Moving Closer to Revenue

Our 2025 survey of 1,200 AI use cases shows a pivot in AI adoption, with leading use cases now centered on **CRM automation, sales enablement, forecasting and lead capture**. Enterprises are steering AI into these areas because they are closest to revenue generating activities, offering measurable links to pipeline growth, margin improvement and customer acquisition. AI as it is being used here is not directly creating new revenue, but it is being **deployed in functions that support and enable growth**, such as faster proposal generation, cleaner CRM data and more accurate forecasts.

This marks a break from 2024, when investment was dominated by chatbots, IT testing and content generation, which were all use cases focused largely on efficiency and service quality. The shift reflects several dynamics: CRM data is structured and revenue-attributable, **integrations with sales platforms are more mature**, and organizations are increasingly **prioritizing areas where ROI can be quantified** in business terms.

Still, results are uneven. As mentioned earlier regarding AI adoption and performance to date, revenue growth remains one of the most underperforming outcomes in AI adoption. This highlights that while enterprises are pointing AI at sales and customer functions, the direct payoff is not yet there.

This trend underscores an important lesson: AI is evolving from cost-saving experiments toward growth-oriented strategies, but the business impact depends on how well organizations integrate data, processes and governance into these high-value workflows.

Source: ISG Market Lens, State of AI, 2025; n=1,200 use cases, not all responses listed

Where are AI Use Cases Being Deployed?



IT Productivity Gains vs. Operational Depth

AI adoption in IT highlights the divide between high-volume use cases and specialist ones. **Copilots and chatbots** dominate the portfolio, but the majority are still in pilot or have achieved a limited rollout. Over two-thirds of copilots remain in testing or early-stage deployment, underscoring the difficulty of scaling front-line productivity tools. These tools attract early attention because they slot easily into developer and support workflows, promising productivity gains without re-architecting infrastructure. Yet the path to full production remains slow, reflecting the sector-wide challenge of turning experimentation into operational scale.

In contrast, use cases such as **predictive analytics model deployment and anomaly detection** represent a smaller share of overall initiatives but show much higher rates of reaching in-production status. Nearly half of the use cases here are live, underscoring that when enterprises do pursue advanced analytics or monitoring, they tend to commit fully.


This reflects the higher stakes: predictive models and anomaly detection often underpin core IT reliability, cybersecurity, or infrastructure cost management. Once validated, they are treated as essential controls rather than discretionary pilots.

Code generation and quality review sits in between, likely because it is a visible but technically demanding area that requires integration into continuous improvement/continuous development (CI/CD) pipelines.

The IT picture is bifurcated: copilots and chatbots attract broad but shallow uptake, while predictive and anomaly initiatives are fewer but more operationalized.

Leading IT AI Use Cases by Status and Investment

Source: ISG Market Lens, State of AI, 2025; n=175 use cases, not all responses listed



Back and Middle Office: Balancing Risk and Efficiency in Core Operations

In the back and middle office, AI adoption is aimed squarely at improving efficiency and compliance. Survey data shows that **legal and compliance use cases** lead the pack, but most remain in pilot phases. This is not surprising: enterprises view compliance as both a high-risk and a high-reward target for AI, so these use cases require extensive validation before scaling. The incentive is strong – automating document review, filings or regulatory monitoring can materially reduce risk exposure – but so are the guardrails.

Forecasting and budgeting is another priority, reflecting pressure on finance teams to accelerate cycles and improve accuracy. Here, adoption mirrors IT’s predictive efforts: relatively fewer use cases, but those pursued are more often in production. Enterprises that prove value in scenario modeling or variance analysis tend to operationalize quickly, since budgeting processes are recurring and well-structured.

By contrast, **AI chatbots and AI-assisted messaging**, which are often aimed at employee services or HR inquiries, attract broad experimentation but limited production deployment. These use cases’ return on investment (ROI) remains harder to demonstrate compared with compliance or finance automation.

Finally, **HR and talent planning** shows the least traction, with few initiatives reaching production. While AI holds promise for workforce forecasting and skills matching, enterprises remain cautious about algorithmic bias and governance.

Source: ISG Market Lens, State of AI, 2025; n=368 use cases, not all responses listed

Legal & Compliance			
% of Use Cases	6%	% in Production	43%
		Investment	\$1.9M
AI Chatbots			
% of Use Cases	5%	% in Production	35%
		Investment	\$0.9M
Forecasting & Budgeting			
% of Use Cases	5%	% in Production	45%
		Investment	\$1.0M
AI Assisted Messaging			
% of Use Cases	5%	% in Production	33%
		Investment	\$2.3M
HR & Talent Planning			
% of Use Cases	4%	% in Production	13%
		Investment	\$0.4M
Supplier Management			
% of Use Cases	4%	% in Production	8%
		Investment	\$2.6M
AI Copilots			
% of Use Cases	3%	% in Production	17%
		Investment	\$2.4M
Supplier Risk Assessment & Monitoring			
% of Use Cases	3%	% in Production	58%
		Investment	\$2.0M
Intranet Search Enhancement			
% of Use Cases	3%	% in Production	40%
		Investment	\$12.6M

Source: ISG Market Lens, State of AI, 2025; n=368 use cases, not all responses listed

Front Office: Experimentation at the Customer Edge

In the front office – which ISG defines as sales, marketing and customer engagement – AI adoption showcases both the promise and the pitfalls of generative AI tools. **AI chatbots** lead all use cases, with a relatively higher share in production than other functions. This reflects years of experimentation with virtual assistants and customer service automation. However, performance remains uneven: while response times and customer satisfaction have improved, ROI in revenue growth and cost savings often falls short of expectations.

AI copilots are gaining traction in sales and marketing teams, assisting with content drafting, pitch preparation and customer research. Yet, like IT copilots, most use cases remain in pilot mode, underscoring the challenges of embedding copilots into frontline workflows where accuracy and trust are paramount.

Pricing optimization and email content generation show meaningful but cautious uptake. These are classic “augmented marketing” use cases: targeted, measurable and repeatable. Enterprises are more likely to scale when they can directly link AI-driven campaigns to conversion metrics.

By contrast, **sales opportunity identification and scoring** is still a nascent use case. Despite clear potential, organizations hesitate to automate decisions about prospect prioritization given the risks of bias and model drift.

AI Chatbots	% of Use Cases 9%	% in Production 49%	Investment \$0.8M
AI Copilots	% of Use Cases 6%	% in Production 24%	Investment \$3.6M
Pricing Optimization	% of Use Cases 6%	% in Production 36%	Investment \$2.4M
Email Content Generation	% of Use Cases 6%	% in Production 41%	Investment \$0.5M
Sales Opportunity Identification/Scoring	% of Use Cases 5%	% in Production 19%	Investment \$2.9M
Customer Segmentation/Journey Mapping	% of Use Cases 4%	% in Production 22%	Investment \$3.0M
Forecasting & Budgeting	% of Use Cases 4%	% in Production 36%	Investment \$3.4M
Personalization	% of Use Cases 4%	% in Production 41%	Investment \$3.6M
CRM Data Entry Automation	% of Use Cases 4%	% in Production 11%	Investment \$4.1M

Leading Use Cases by Industry

Industry	Top Use Case
Banking & Financial Services	AI Chatbots - Lead Capture
Entertainment & Media	AI-Assisted Script & Content Development
Healthcare	Voice-to-Text Clinical Note Generation
Information & Services	Virtual Developer Agents / AI Chatbots
Insurance	Pricing Optimization
Life Sciences	Predictive BI Dashboards
Manufacturing	Market Demand Prediction
Oil & Gas	AI Copilots - Infrastructure Administrators
Retail	Sales Forecasting
Telecommunications	AI Copilots - Product Managers
Travel, Tourism & Hospitality	AI Chatbots - Internal Communications
Utilities	Predictive BI Dashboards

Source: ISG Market Lens, State of AI, 2025; n=1200 use cases, not all responses listed


Industry-Specific Use Cases: Variation Emerges

Horizontal use cases like chatbots, copilots and forecasting remain dominant across industries, reflecting their broad applicability and relative ease of deployment. Yet the picture changes when looking at industry-specific priorities.

For example, the healthcare industry emphasizes clinical support such as clinical note generation. This differs from the manufacturing and retail industries, which are focused on forecasting use cases, and the banking industry, which seems to be pursuing lead capture and email content generation to support customer acquisition.

These vertical-specific use cases are fewer in number but often more strategic, requiring specialized data and domain knowledge. Their emergence highlights a growing shift: while horizontal tools drive scale and efficiency, industry-specific applications define competitive differentiation and reveal how AI adoption is beginning to diverge more sharply sector by sector.

The motivations and metrics used to measure success also vary. For example, most enterprises target efficiency and productivity outcomes from AI. But retail has taken a different path, seeking to drive sales and revenue growth and enhance customer experience.





Executive Summary

The Evolving Value of Applied AI

State of the Market

The Enterprise AI Ecosystem

Models Meet Markets

Requirements

Key Contributors

The Enterprise AI Ecosystem

The Enterprise AI Ecosystem

Applied AI is not a single product but a system of interdependent parts. To achieve business value, enterprises need to assemble and coordinate multiple layers of technology and expertise: the platforms that power model development, the software that embeds AI into business functions and the services that connect and maintain them. What matters is not the presence of any one layer but how effectively they work together.

AI platforms provide the strategic foundation. They support the full lifecycle needed to prepare, train, deploy and govern models; they bridge the gap between experimentation and scaled adoption. While enterprise software applications accelerate time to value by embedding AI into front- and back-office processes such as customer service, HR or finance, they also bring their own trade-offs around differentiation and control. Service providers complement both differentiation and control challenges by delivering the integration, modernization and managed services needed to keep platforms and applications aligned with business outcomes.

Together, these elements form an ecosystem that is far more than the sum of its parts. Enterprises that recognize and manage this complexity are better positioned to scale AI responsibly, avoid the pitfalls of fragmented investments and turn technological potential into operational results.

AI Capability Reference Architecture Design

AI Consumption

Enterprise Apps • Sandbox Area • User Interface • Agents / APIs

AI Infusion for Enterprise Applications

Model Integration via API Management • Platform Compatibility for Model Deployment • Automated Model Integration • Orchestration, Memory & Reasoning

AI Model Development / Deployment

Model Training & Tuning • Established Model Deployment • Model Performance Monitoring • Model Lifecycle Management

AI-Enabled Data

Data Ingestion • Data Transformation Engineering • Data Access & Delivery • Contextualization & Augmentation • Data Quality, Validation & Governance • Data Platform

AI-Ready Infrastructure

Model-Aware Computing • AI-Optimized Storage • Domain-Specific Accelerator • Intelligent Edge • Networking & Communication • Hybrid Cloud • End-User Compute (EUC)

AI Management

- AI Systems Management
- Secure AI
- AI FinOps
- Responsible AI
- Compliant AI

AI Tools

- Development & Deployment Tools
- Automated Machine Learning Platforms
- AI Asset Management Tools
- Collaborative Platforms for AI Development
- Model Explainability & Interoperability Tools

AI Platforms

The Strategic Foundation of Enterprise AI

AI platforms, along with data platforms, are emerging as the core infrastructure for enterprise AI adoption. They are integrated environments that enable organizations to prepare, train, deploy and maintain AI models. Though these platforms have a longer history, only now are they becoming central to how enterprises operationalize AI at scale.

Modern AI platforms support the entire model lifecycle: preparing data, training and optimizing models, deploying models into enterprise processes for inferencing and then maintaining the models through governance, monitoring and retraining. This holistic view is essential as enterprise use of AI moves from isolated pilots to mission-critical applications that require accuracy, resilience and compliance.

Though AI platforms have existed for decades, their adoption has long been impeded by two barriers. Historically, processing the large datasets required for accurate modeling was prohibitively expensive. Advances in scale-out computing and object storage have now made this feasible. Second, recruiting and retaining data science talent has been – and remains – one of the top challenges.

Recent breakthroughs in GenAI have further democratized access, extending AI use beyond data science teams. With agentic AI, which is capable of autonomous actions, the demand for robust platforms has intensified even further.

Data Is Essential to AI

AI platforms need clean, well-organized and regulatory compliant data. Challenges around data readiness are addressed in detail later in the report, where we examine how legacy data models are breaking under AI's demands and why a new “data by design” approach is now essential.

The Governance Gap

Despite these advances, AI governance remains a weak point. Platforms are beginning to add guardrails, explainability and compliance monitoring, but many of these capabilities are not yet fully reliable. Risks like bias, drift and compliance breaches could erode trust and slow adoption.

This governance gap is not unique to AI platforms. As explored later in this report, enterprises are finding that legacy governance frameworks are cracking under pressure in the AI era. While platform guardrails are emerging, a full resolution requires rethinking governance models, operating structures and data practices – topics we return to in detail in the *Foundational Requirements for Achieving Value* section.

What's Next

AI platforms are becoming strategic foundations for enterprise AI. They bridge experimentation and production, and embed governance, compliance and operations. As such, they are poised to serve as the operating system of applied AI, enabling enterprises to scale AI responsibly and competitively.

Enterprise Software

As enterprises embark on their AI journeys, an important consideration is the build-versus-buy decision. Time to value can often be accelerated with pre-built software applications. Software application providers have been investing heavily in AI-based capabilities across customer experience, human capital management, customer relationship management, finance and supply chain applications. For example, contact center applications now include call transcription and summarization, and recruiting applications offer candidate matching. These built-in capabilities may offer enterprises a faster, more reliable path to value than building custom solutions.

AI Capabilities in Enterprise Software

Customer Experience	HCM	CRM & Sales	Finance	Supply Chain
Transcription & summarization	Employee self-service	Lead / opportunity scoring	Forecasting using predictive analytics	Predictive analytics - demand forecast
Self-service chatbots	Candidate matching	Revenue forecasting	Outlier & error detection	Real-time monitoring
Agent guidance (next-best-action)	Inferred skills	Activity collection & summarization	Analytics and reporting automation	S&OP (limited)
Agent quality assessments	Skills alignment	Account opportunity generation	Automated flux analysis	Concurrent planning
Agent self-management	Onboarding	Recommendation engine	Accounting matching & reconciliations	Logistics forecasting
Call risk scoring for compliance or fraud	Digital coach	Offers & pricing	Automated invoice ingestion & posting	Reporting automation
Workload forecasting	Compensation analysis	Upset / cross-sell	Continuous assurance (fraud)	Process management
Auto Pill redaction	Payroll trend analysis	Data cleansing	Process management	Scenario planning
"Knowledge management"	Skills assessment & insights	Multi-language text descriptions		Route optimization
	Just-in-time performance coaching			

Advantages of Buying

Talent and systems are the biggest barriers or delays in moving to production. Nearly two-thirds of enterprises cite talent as the biggest barrier; 57% cite systems as the biggest barrier. Enterprises can overcome both of these issues by acquiring and deploying enterprise software applications. Addressing these two issues can help accelerate the time to value by making AI-based capabilities available immediately. Built-in guardrails also can help mitigate governance concerns that impede custom AI development efforts.

In addition, software providers take on the responsibility for maintaining and enhancing AI capabilities, which frees up internal resources to focus on the business processes they are intended to support. Also, software providers may have a broader set of data from across their customer base on which to base AI models, resulting in more accurate predictions and processes.

Disadvantages of Buying

However, it doesn't make sense to put the "AI cart" before the horse – the horse in this case being the application itself. Enterprises need to evaluate the functionality of the application separate from any AI capabilities it may have to ensure it meets their requirements. Only then should they consider the AI capabilities. When evaluating vendor-provided AI capabilities, it is important to consider how they will help an enterprise differentiate itself from others. If every company has the same capabilities, where is the competitive advantage? If a combination of pre-built capabilities and custom capabilities is necessary, then an enterprise buyer needs to ensure the application can support customization.

Service Providers

Putting the Applied in “Applied AI”

Service providers play an important role helping enterprises turn the promise of applied AI into operational reality. They provide two complementary sets of capabilities: 1) custom services for AI and analytics, and 2) platform-centric services for implementing and managing AI ecosystems. Together, these services accelerate adoption, ensure scalability and align technology with business outcomes.

Custom AI and Analytics Services

Custom services tailor AI to enterprise-specific needs. Providers not only advise on use-case selection, ROI potential and governance, they also execute through data engineering, model development and workflow integration supported by machine learning operations (MLOps) and large language model operations (LLMOps) for continuous retraining and monitoring.

Data modernization is key. Providers can migrate legacy systems to cloud-native environments, create trusted data foundations and connect them to business outcomes. Increasingly, GenAI is layered onto analytics to support data storytelling and decision automation.

AI Platform–Centric Services

Platform services ensure enterprises maximize investments in AI platforms. These services include advisory offerings, which help enterprises align platform choices with their business strategy and compliance needs. These services also include AI platform implementation support, which involves configuring data pipelines, integrating across multicloud environments and embedding AI workflows within platforms.

Purchasing AI platform-centric services in a managed services model can help enterprises sustain the performance and cost efficiency of AI platforms post-deployment. These services are designed to deliver monitoring, automation and governance controls that keep platforms secure, resilient and compliant.

Convergence of Custom and Platform Services

Although distinct, custom services and platform services are tightly linked. Custom solutions depend on robust platforms, and platforms deliver value only when tied to business-specific AI applications. Providers increasingly offer integrated models that combine both and help enterprises put into place responsible AI practices such as fairness, explainability and compliance throughout.

A notable example of this integrated model is the rise of data tower-as-a-service. By blending technical execution with business context, providers are building and operating dedicated data environments that accelerate AI maturity without waiting for internal alignment. This approach not only speeds readiness but also addresses the acute shortage of data engineering and AI operations talent.



AI Pricing: Models Meet Markets

AI Pricing: Models Meet Markets

Pricing is where the promise of AI collides with enterprise reality. In 2025, two clear patterns have emerged.

On one side, enterprises are buying **agents as software** – paying for the code, access to the model(s) and the platform – then taking responsibility for running and governing them. This feels familiar, echoing the shift to Software-as-a-Service (SaaS) and cloud. But as with cloud's early years, hidden costs can accumulate quickly. Beyond the license fee, buyers must also account for:

- **IP creation costs** to build, fine-tune or domain-train the agent.
- **Ongoing maintenance costs** for monitoring, tuning and updates.
- **Token consumption** across inference and multi-model workflows.
- **Infrastructure costs** to host and scale the agent.

Without a **financial operations-style** (FinOps-style) **discipline for agents**, organizations risk repeating the cloud sprawl experience: growing value alongside unpredictable bills.

On the other side sit **transaction-centric contracts** – the traditional business process outsourcing (BPO) or managed service deals priced per invoice, per claim or per ticket. Agents are reshaping how these services are delivered. Work once performed by people or scripts is increasingly handled by agents that suggest, execute or orchestrate. The result is tension: familiar unit prices are applied to unfamiliar execution models. Providers are experimenting, but many models obscure value more than they reveal it.

Grounding with Seven Pricing Principles

No matter how contracts evolve, ISG believes any viable model must hold to these seven principles:

- 1. Fairness:**
Value exchange must be equitable for both parties.
- 2. Variability:**
Charges flex with consumption or measurable outcomes.
- 3. Consistency:**
Pricing models remain stable and comparable across providers.
- 4. Competitiveness:**
Prices can be benchmarked against the market.
- 5. Incentives:**
Contracts reward efficiency and innovation.
- 6. Risk Awareness:**
Allocation of risk is transparent, especially as autonomy grows.
- 7. Proportionality:**
Charges map to true cost drivers, not proxies.

Pricing models that stray from these principles – whether in a digital full-time employee (FTE) model, flat output rates or raw token billing – create friction, disputes and adoption stalls.

Autonomy-Level Pricing: Aligning Principles with Practice

The most promising bridge for today’s transaction-based services is **autonomy-level pricing (ALP)**. ALP doesn’t discard the familiar **resource unit (RU)**, but augments it by tagging each unit with the autonomy level used to deliver it.

Autonomy Levels Defined	
Autonomy Level	Description
AL0	Fully manual execution; no AI involved.
AL1	AI suggests, human executes.
AL2	AI executes, human validates.
AL3	AI executes, human audits exceptions.
AL4	Fully autonomous execution, governed only by policies and audits.

With ALP, a claim processed at AL1 is priced differently than one at AL3. The construct keeps procurement logic intact – “per invoice” is still a valid unit – but adds transparency into how it was delivered.

Example: Finance Record-to-Report			
Process	ALP Level	Execution Behavior	Governance Role
Journal Entry Creation	AL1	AI drafts; human posts	Human-in-the-loop (HITL)
Accrual Calculations	AL2	AI proposes; manager verifies	Verification required
Account Reconciliation	AL2	AI matches; human confirms	HITL required
Variance Analysis	AL3	AI surfaces anomalies	Audit on exception
Close Checklist Management	AL3	AI orchestrates tasks	Escalation-based
High-Volume General Ledger Posting	AL4	AI posts within thresholds	Retrospective audit

This mapping shows how ALP enables nuanced pricing. At AL1 and AL2, human involvement remains central, so charges reflect collaborative effort and validation. At AL3, audit-on-exception makes autonomy viable without removing oversight. At AL4, full autonomy is priced lower per unit, but subject to auditability and fallback obligations.

Why ALP Works

- **Fairness:** Buyers see cost reductions as autonomy rises; providers retain margin by not overcommitting upfront.
- **Variability:** Prices flex with real execution behavior, not assumptions.
- **Consistency:** Contracts preserve the RU anchor, making deals comparable across towers and providers.
- **Competitiveness:** ALP levels can be benchmarked against industry peers.
- **Incentives:** Providers are rewarded for safely advancing autonomy; buyers capture savings without waiting for maturity.
- **Risk Awareness:** Each ALP tier defines governance roles, making risk explicit.
- **Proportionality:** Charges map to autonomy, service-level agreement (SLA) burden and error risk – not arbitrary token counts.

In short, ALP operationalizes the seven pricing principles rather than bypassing them.

ALP and the Evolution of RUs

Resource units (RUs) are the backbone of enterprise finance. They define budgets, enable benchmarking and anchor vendor scorecards. But they are **agnostic to how the work is done**: a “ticket” costs the same whether resolved by a human analyst or an AI agent.

ALP makes RUs smarter by adding **execution context**. A ticket may be billed at AL1, AL3, or AL4 depending on how it was resolved. An invoice may be tagged AL2 if the AI executed the work, but a human validated it. This simple overlay transforms static units into **intelligence-tagged units**, giving sourcing teams transparency without disrupting familiar models.

Benefits include:

- Transparent alignment of price with execution maturity.
- Credit to providers for delivering higher autonomy safely.
- Benchmarkable comparisons across suppliers.
- Audit trails that satisfy governance and regulatory needs.

Importantly, ALP is not permanent scaffolding. Just as early BPO contracts required **re-basing** when automation matured, ALP requires regular recalibration as the agentic frontier continues to move. Multipliers tied to yesterday’s autonomy levels should not harden into tomorrow’s outdated labor proxies.

Where Pricing Stands Today

The state of AI pricing is unsettled. **Software-style deals** are stabilizing but require new cost visibility. **Transaction-centric deals** remain messy, but ALP offers a governance-aware bridge. The models will keep evolving. The principles won’t. Enterprises that hold to them – and use ALP to modernize RUs – will turn AI’s potential into predictable, governed value.

Foundational Requirements to Achieving AI Value

Foundational Requirements to Achieving AI Value

AI is only as powerful as the foundations on which it is built. While enterprises have poured resources into pilots, proofs of concept, and high-profile generative and agentic AI initiatives, as mentioned earlier in this report, results remain uneven.

Organizations are learning that without strong data, a clear approach to adoption and rigorous governance, the promise of AI cannot be fully realized. In fact, many of today's most visible failures in scaling AI – from inaccurate copilots to stalled agent rollouts – stem from brittle foundations rather than shortcomings in the technology itself.

Three core requirements consistently emerge as the bedrock for AI at scale: data by design, adoption accelerators and governance. Together, they form a framework for moving from experimentation to sustained enterprise value.

Data by Design: Building a Foundation for AI at Scale

From Legacy Systems to AI-Ready Data

Traditional enterprise data systems were optimized for reporting and compliance, not for the dynamic, multimodal

demands of AI. They were designed around batch pipelines, rigid schemas and governance models that favored control over adaptability.

Generative and agentic AI have upended this paradigm. Instead of asking for curated datasets weeks in advance, these systems thrive on real-time, unstructured and heterogeneous data streams – whether drawn from transactional systems, IoT sensors or customer interactions. Of course, providing good context is still important to improve generative and agentic AI performance.

This upended paradigm has exposed an enormous gap in readiness: while AI needs contextually rich and faster data, most enterprises are still encumbered by fragmented landscapes of thousands of applications, siloed semantic definitions and legacy architectures that prevent data from being contextualized and operationalized.

AI's evolution has revealed what many already suspected: data is less an asset to be hoarded and more a liability unless it can be transformed into usable, trusted and timely intelligence.




A Three-Mode Operating Model

To address this challenge, enterprises are shifting from the idea of a dual-speed operating model to a more nuanced three-mode cycle as the operating model. Think about the modes in this way:

- **Optimization mode** focuses on stable, proven-value data with a direct link to ROI – the financial reconciliations, regulatory reports and risk models where precision, governance and repeatability are paramount. These assets justify heavy investment in pipelines, controls and architecture because their value is already established.
- **Innovation mode** supports exploratory use cases, pilots and AI-driven experiments that demand speed and flexibility. Here, governance is intentionally lighter, architectures are often bypassed and the goal is discovery via testing new sources, models and unstructured or multimodal data.
- **Adoption mode** serves as the bridge between the two. It is not simply “deploying a pilot,” but rather it is capturing the governance, engineering and architectural lessons learned in innovation and applying them across the enterprise. If a proof of concept required new governance for unstructured data, adoption mode ensures that governance becomes part of the enterprise’s standard operating model.

Enterprises that neglect adoption mode risk falling into two traps: years-long “boil the ocean” data programs that never deliver value or fragmented proofs of concept that cannot scale. By treating adoption as a distinct discipline, organizations create a structured pathway from innovation to optimization – ensuring experiments do not remain isolated, and optimizations are continually refreshed by new ideas.

From Innovation to Optimization: Evolving the Operating Model



Semantic Clarity and the Rise of the Data Tower

One of the most significant barriers to AI value remains semantic inconsistency. Without shared definitions, even structured data loses value. Enterprises are beginning to respond by elevating data operations into first-class business functions – what ISG describes as **the data tower**. These centralized yet business-aligned constructs emphasize semantic clarity, domain-specific data products and accountability through key performance indicators (KPIs).

Application-Native Data Governance

Perhaps the most transformative shift in data design for AI is the movement toward application-native governance. Historically, enterprises pulled data from source applications into curated data lakes and warehouses before it was deemed safe for consumption. Agentic AI challenges this sequence.

By acting directly on back-end application data in real time, agents take the shortest path in comparison to traditional pipelines. For organizations to maintain control, governance will start to be enforced at the source, with validation, quality scoring and even write-back mechanisms embedded directly in the systems of record. This will create a significant challenge for enterprises with a large legacy footprint.

Preparing for this future today – where governance is dynamic, embedded and close to the point of data creation – is becoming a competitive advantage. Enterprises that continue to rely solely on centralized batch models risk creating bottlenecks that agents will simply work around.

Data by design means shifting from static curation of data to dynamic orchestration across optimization, innovation and adoption cycles. It is about building semantic clarity, elevating data into a business function and embedding governance directly where the data lives and is actively used. Without this foundation, AI will remain a patchwork of pilots rather than a platform for transformation.

Adoption Accelerators: From Experimentation to Scale

Learning From Past Failures

Many enterprises remain stuck in what ISG calls the pilot-to-production dilemma. Early AI pilots, often celebrated for their innovation, too rarely progress to enterprise-grade deployments. Two delivery models have dominated, and both have shortcomings:

- **Boil the Ocean:** Large, multi-year data transformation programs promise to “get our data right” before tackling AI. These efforts often stall under their own weight with governance debates consuming more energy than innovation.
- **Bypass the Mess:** In contrast, frustrated leaders often spin up siloed pipelines for immediate needs, creating shadow IT and reinforcing fragmentation.

Neither extreme works. Enterprises should instead operate in a cycle of innovation, adoption and optimization – experimenting rapidly, codifying lessons through adoption, and hardening them into scalable, compliant processes. This cycle prevents years-long delays while ensuring experiments don’t remain disconnected from broader enterprise value.

The Role of Organizational Change

Adoption challenges are not purely technical. In fact, organizational readiness is often the bigger constraint. Agentic AI in particular forces companies to rethink workflows, decision rights and even job roles beyond IT. Human roles become elevated to focus on enriching semantics and knowledge; managers become orchestrators of outcomes rather than just task supervisors; and analysts evolve into curators of context rather than manual data processors.

Change management is not optional. Just as with the digital transformation era, success depends on aligning incentives, redesigning processes and fostering a culture of trust in autonomous systems. Enterprises that treat AI adoption as a pure IT project – without investing in training, communication and change management – consistently report underwhelming ROI.

Talent and Provider Ecosystems

The scarcity of AI and data talent compounds the adoption challenge. Few enterprises have enough professionals who understand both their complex data environments and AI tooling.

Most rely on service providers for augmentation, yet even here the choice of partner is critical. Providers have become adept at designing commercial models that align to outcomes, blending efficiency with flexibility. More mature providers offer cross-functional, outcome-based delivery models that combine domain expertise, data science and transformation skills.

This reliance on providers reflects a broader truth: adoption is not just about tools, but ecosystems. The enterprises accelerating the fastest are those that leverage platform partners (AI platforms, data clouds and/or hyperscalers), SaaS vendors who are embedding AI, and service providers with deep contextual expertise. Together, this ecosystem helps organizations scale AI beyond pilots and integrate it into the workflows and processes that matter most.

Avoiding the Value Gap

Despite heavy investment, many enterprises report an AI value gap: efficiency gains materialize, but business growth impacts lag. In ISG surveys, nearly half of enterprises expect meaningful AI-driven growth only in 2026 or beyond. To accelerate adoption, leaders should focus not only on cost savings and automation but also on growth-oriented use cases – new product design, customer experience and industry-specific reinvention.

Governance: Guardrails for Responsible AI

From Control to Dynamic Oversight

Traditional governance frameworks emphasized static controls: data stewardship committees, manual reviews and policies enforced after the fact. Such models are simply too slow for AI that learns, generates and acts in real time. In the agentic era, governance will need to evolve into dynamic oversight: real-time monitoring, embedded escalation logic and adaptive feedback loops.

This means governance must operate not only on the data and models but also on the **actions AI takes**. Agentic AI introduces a new class of risks: flawed purchasing decisions, inappropriate customer interactions or operational disruptions. Governing action – not just output – becomes essential.

Ethical and Legal Implications

As enterprises deploy AI deeper into their value chains, they face heightened ethical and legal scrutiny. Key risks include:

- Hallucination and inaccuracy, leading to poor or unsafe decisions. Bias and inclusivity gaps, perpetuating inequalities.
- Privacy and intellectual property (IP) leakage, as models consume and generate sensitive data.
- The need to maintain autonomy for data owners and custodians. An enterprise must be able to revoke permissions to a data set at any time, and when this happens, models should be re-trained without the revoked dataset in a timely manner.
- Regulatory non-compliance, as global frameworks like the European Union AI Act and the Digital Operational Resilience Act (DORA) for banking, financial services and insurance (BFSI) firms set strict requirements.

Organizations are responding by embedding responsible AI principles into their frameworks.


Leading providers are building ethical AI pillars, while enterprises are adopting policy-as-code, lineage tracking and trust scoring to automate compliance. Transparency and explainability are becoming table stakes: employees, customers and regulators alike demand to understand how AI is making decisions.

Governance as a Differentiator

While governance is often seen as a constraint (or a gap, in terms of software), it **increasingly functions as a differentiator**. Enterprises that can demonstrate robust, responsible AI practices gain trust with regulators, customers and employees. Providers, too, are competing on their ability to deliver **industry-specific governance frameworks**, from risk-aware models in financial services to compliance-heavy controls in healthcare.

The most forward-looking enterprises are reframing governance not as an overhead function but as a value enabler. By embedding compliance, ethics and observability into their architectures from the start, they can scale AI faster, avoid costly rework and position themselves as trusted leaders in their industries.

Governance for AI is shifting from static control to dynamic, embedded guardrails that protect trust, enable speed and differentiate leaders from laggards.



What's Next

The story of AI in the enterprise is no longer about whether the technology works. It is about whether organizations are ready to adopt the changes and processes which have been shown to lead to success. Without data by design, AI remains starved of the reliable, contextual inputs it needs. Without accelerated adoption, pilots never scale into transformative workflows. Without governance, trust collapses, and risk outweighs reward.


The enterprises that win in the AI era will not be those with the flashiest pilots, but those that embed foundations in their culture and treat it as strategy. They will build three-mode data operating models, orchestrate adoption with ecosystems and change management, and embed governance directly into their AI workflows. In doing so, they will not only avoid the pitfalls of hype cycles but will position themselves to capture the durable value of intelligence at scale.

Key Contributors




Loren Absher

Director & Americas Lead - Applied AI Advisory




Alex Bakker

Distinguished Analyst




Dorotea Baljevic

Director & ISG EMEA AI Lead




Steve Hall

President, ISG EMEA, and Chief AI Officer



Stanton Jones

Distinguished Analyst



Olga Kupriyanova

Director - AI & Data Engineering



Dave Menninger

Executive Director - ISG Software Research & Distinguished Analyst

With Special Thanks To

Taylor Card


Gowtham Sampath


Michael Dornan

Hannah Stapleton

Related ISG Research

- ISG State of the Agentic AI Market Report
- ISG Buyers Guides™ for AI Platforms
- ISG Buyers Guides™ for Data Platforms
- ISG Provider Lens™ for Advanced Analytics & AI Services
- ISG Provider Lens™ for Generative AI
- ISG Index Insider for complimentary weekly briefings





ISG (Nasdaq: IHI) is a global AI-centered technology research and advisory firm. A trusted partner to more than 900 clients, including 75 of the world's top 100 enterprises, ISG is a long-time leader in technology and business services that is now at the forefront of leveraging AI to help organization achieve operational excellence and faster growth. The firm, founded in 2006, is known for its proprietary market data, in-depth knowledge of provider ecosystems, and the expertise of its 1,600 professionals worldwide working together to help clients maximize the value of their technology investments.