

Generative AI Services — Large and Midsize

A comparative analysis of service providers
focused on unlocking and maximizing the
potential of GenAI

QUADRANT REPORT | OCTOBER 2026 | GLOBAL



Customized report courtesy of:



Executive Summary 03

Provider Positioning 08

Introduction

Definition	14
Scope of Report	15
Provider Classifications	16

Appendix

Methodology & Team	44
Author & Editor Biographies	45
About Our Company & Research	47

Strategy and Consulting Services — Large 17 – 22

Who Should Read This Section	18
Quadrant	19
Definition & Eligibility Criteria	20
Observations	21

Strategy and Consulting Services — Midsize 23 – 29

Who Should Read This Section	24
Quadrant	25
Definition & Eligibility Criteria	26
Observations	27
Provider Profile	29

Development and Deployment Services — Large 30 – 35

Who Should Read This Section	31
Quadrant	32
Definition & Eligibility Criteria	33
Observations	34

Development and Deployment Services — Midsize 36 – 42

Who Should Read This Section	37
Quadrant	38
Definition & Eligibility Criteria	39
Observations	40
Provider Profile	42

*Report Author: Gowtham Kumar
Sampath and Sushree Acharya*

GenAI differentiation is shifting from AI capability to repeatable enterprise change

This ISG Provider Lens® study evaluates large and midsize providers delivering GenAI services globally, spanning strategy and consulting, development, deployment, integration, governance and production-scale operations. The assessment examines how providers help enterprises move beyond experimentation toward secure, contextualized and economically sustainable GenAI, translating rapidly evolving model capabilities into enterprise workflows and measurable business outcomes.

Market Context

GenAI enters a more commercially disciplined phase

GenAI investment remains resilient even as enterprises operate under tighter technology spending conditions. ISG Market Lens™ 2026 IT Budgets and Spending Study finds average

overall IT budget growth tightening to 1.2 percent, from 1.8 percent in 2025, while AI and security continue to receive comparatively stronger investment. The 2Q26 ISG Index reinforces this divergence: demand for AI-driven infrastructure and software is expanding rapidly, while the economics of managed services remain under greater pressure. This divergence is sharpening the focus on AI economics, as enterprises seek measurable productivity gains and returns, while providers face growing pressure to translate AI-enabled delivery into sustainable service economics.

As GenAI absorbs a larger share of enterprise AI investment, expectations are shifting beyond model capabilities and successful deployment to assessing whether GenAI can remain reliable, governed and economically viable in production. ISG Research finds that nearly half of enterprise AI budgets are allocated to GenAI, while the broader AI platform market is expanding around lifecycle development, deployment, governance and optimization.

Regulation, sovereignty and skills add to this operating challenge. Governance increasingly extends across the AI lifecycle

As **GenAI** becomes abundant, **differentiation** shifts to **enterprise context, work and outcomes.**



Executive Summary

as models, enterprise information and regulatory requirements evolve, while sovereignty is expanding from data locality toward architectural autonomy and choice across infrastructure, models, software and deployment environments. This creates tension between greater enterprise control and the complexity of maintaining heterogeneous environments. AI expertise also remains constrained, increasing the importance of engineering, automation and reusable assets that allow GenAI to scale without specialist resources growing proportionately.

Enterprise Priorities

Prove value while controlling the economics of GenAI

Enterprise expectations are moving beyond access to GenAI. ISG Market Lens™ 2026 AI Value Study shows organizations emphasizing staff productivity, agility, revenue and profit growth and direct cost reduction, while task automation, process optimization and data analysis have generated some of the most tangible value recently. Enterprises increasingly evaluate initiatives through predefined KPIs,

productivity and ROI rather than adoption or pilot counts alone.

Implementation evidence helps explain why productivity remains prominent, as GenAI reduces the effort required to retrieve expertise, interpret information, synthesize evidence, develop software and prepare decisions. Its immediate economic significance often stems from changing the economics of cognition, making information-intensive work faster, less resource-intensive and accessible at scale.

However, productivity should not be interpreted as evidence of limited AI maturity or ambition.

Material business value can emerge from augmentation well before roles or operating models change, and a relatively conventional implementation applied across high-volume work can create greater economic impact than an advanced application with limited operational reach. Enterprises, therefore, need to distinguish between AI presence, AI value and AI dependence rather than treating use-case proliferation or architectural sophistication as proxies for impact.

As usage scales, the economics of operating GenAI also become more important.

Model-token consumption, infrastructure, data preparation and ongoing production requirements can make costs difficult to predict, requiring enterprises to balance performance against inference cost, latency and workload scale rather than focusing primarily on initial model capability.

Make enterprise knowledge usable and trustworthy

Enterprise context becomes increasingly important as GenAI moves deeper into work. Applications are progressing from generation and retrieval toward synthesis, interpretation and contextualization, allowing documents, policies, code, customer interactions, domain knowledge and operational evidence to become reasoning inputs.

Hence, connecting models to information is only part of the requirement when outputs must be accurate, traceable and suitable for downstream use. Enterprises increasingly need to preserve authoritative sources, permissions, provenance, business relationships and

terminology while ensuring that the context available to GenAI remains current as underlying enterprise knowledge evolves.

This makes data and knowledge readiness more directly connected to the value of GenAI. The challenge is no longer simply providing models with more information, but organizing the information required for them to interpret enterprise context appropriately and produce outputs that employees, applications and workflows can consume with confidence.

Move from deployment toward dependable production

Production introduces requirements that were less visible during experimentation because models, enterprise information, usage patterns and operating conditions continue to evolve after deployment. Consequently, evaluation, security, governance, observability, context quality and lifecycle management become ongoing operating requirements rather than one-time implementation activities.

Enterprises also need to determine where probabilistic outputs are acceptable and



where business processes require stronger validation, deterministic controls or human judgment. Structured, validated and traceable GenAI outputs can provide an important bridge between individual productivity and process value, allowing machine cognition to participate in operational work without requiring AI itself to own execution.

Preserve architectural flexibility without creating unnecessary complexity

Different workloads create different requirements for quality, cost, latency, privacy, sovereignty and deployment, making a single model or platform increasingly unlikely to provide the optimal answer for every enterprise requirement. Enterprises, therefore, need sufficient architectural flexibility to benefit from ongoing model innovation and changing economics without repeatedly rebuilding applications around a preferred technology. Greater optionality introduces its own trade-offs. Multi-model environments, heterogeneous infrastructure and sovereign deployment choices can increase operating complexity,

making portability and architectural control valuable only when matched by the engineering and governance capabilities required to manage them.

Determine how work should change

The implementation reality exposes a deeper challenge, as enterprises are embedding GenAI into existing work considerably faster than they are redesigning work around it. GenAI is already improving knowledge work, software development, customer interactions, decision preparation and information-intensive processes, while evidence of equivalent changes to roles, handoffs and operating models remains much thinner.

Enterprises should understand this distinction because GenAI can create significant economic and operational value without fundamentally transforming how an organization operates. The next frontier is therefore not simply expanding adoption but determining where cheaper and more scalable machine cognition justifies changing tasks, handoffs, roles and the design of the underlying work.

Provider Dynamics

Differentiation moves from retrieval toward enterprise context and workflow knowledge

Providers are extending their propositions beyond model implementation and generic retrieval as enterprise requirements become more contextual. Enterprise knowledge, metadata, semantic structures, permissions, knowledge graphs and domain information are increasingly combined with mechanisms for maintaining context as business information changes. Generic RAG is consequently becoming a common foundation, while differentiation comes from organizing the knowledge required for GenAI to interpret and participate in specific enterprise work.

This development is increasingly converging with verticalization. Broad industry expertise remains important, but providers are translating it into more granular subdomain, functional, process and workflow knowledge that can be encoded into reusable assets and contextual structures. Context engineering and vertical specialization, therefore, converge where

providers translate business expertise into the information, relationships and rules required to make generic AI useful for specific enterprise activities.

Responsibility expands from model deployment toward systems engineering and continuous operation

Production implementations increasingly combine GenAI with retrieval, conventional machine learning, optimization, deterministic rules, APIs, enterprise applications, workflow controls and human expertise. The resulting enterprise AI architecture is becoming more compositional, with GenAI complementing traditional AI, automation and deterministic systems while each technology performs the parts of the problem to which it is best suited.

Technical depth consequently increasingly depends on engineering these components as a coherent enterprise system, determining where probabilistic intelligence adds value while retaining precision, optimization, transaction integrity and control where the business requirement demands them.



Executive Summary

This broader engineering responsibility extends to continuous evaluation, observability, tracing, governance, security, lifecycle management and ongoing cost and performance optimization as models, enterprise context, usage and operating conditions continue to evolve after deployment. Production is becoming an operating responsibility rather than a deployment milestone.

In addition, architecture becomes an optimization problem as quality, cost, latency, privacy, sovereignty and infrastructure requirements influence model and deployment choices across workloads. Model evaluation, abstraction, routing and flexible deployment capabilities are becoming more visible in provider portfolios, although implementation evidence suggests that sophisticated enterprise-wide model portfolio management remains less mature than providers' capabilities in this area. The opportunity is therefore credible, but dynamic multi-model optimization should not yet be mistaken for the market center.

Sovereignty and rapid model evolution reinforce the value of modularity and

portability, particularly where enterprises need greater control over infrastructure, models and deployment environments. Contextual components, connectors, evaluation frameworks and workflow assets that can operate within heterogeneous client estates can preserve this choice without requiring wholesale adoption of a proprietary AI stack.

Industrialization moves from reusable technology toward repeatable enterprise implementation

Reusable code, accelerators, reference architectures and AI factories continue to improve delivery efficiency, but these capabilities are increasingly widespread. More advanced industrialization is moving toward reusable domain knowledge, contextual components, workflow patterns, integrations, evaluation mechanisms, controls and operating practices, allowing implementation learning to accumulate beyond the technical components themselves.

Repetition is not repeatability as experience from similar projects demonstrates delivery exposure, while genuine repeatability requires

evidence that knowledge from earlier implementations systematically improves subsequent delivery and reduces the extent to which context, workflow design, integration and assurance must be recreated for every client. The stronger competitive progression, therefore, moves from reusable technology toward reusable implementation knowledge and, ultimately, repeatable business outcomes.

Large providers can draw on extensive integration capabilities, operating experience and exposure to complex enterprise environments, while midsize providers can create comparable leverage in selected areas through focused engineering, specialization and reusable IP. Across both populations, the quantity or sophistication of GenAI technology is becoming a less discriminating measure of leadership as the market increasingly tests whether those capabilities can translate into enterprise outcomes with repeated success.

Outlook

Scarcity moves closer to the enterprise

Over the next 12-24 months, advances in models, retrieval, engineering frameworks and

production tooling will make important layers of GenAI increasingly accessible and repeatable. This will not make enterprise GenAI simple. Instead, scarcity will shift toward capabilities that remain harder to standardize, including maintained enterprise context, composite systems engineering, workflow knowledge, continuous assurance and the ability to translate implementation experience into repeatable business outcomes.

Enterprise maturity will consequently become harder to judge from technology adoption alone. A growing portfolio of GenAI use cases demonstrates penetration but not equivalent value or structural dependence, while faster access to information and decision preparation do not necessarily establish better decisions. The same distinction applies to technical sophistication, because complex architectures can produce limited business impact, while comparatively straightforward implementations operating at scale can materially affect cost, quality or risk.

The competitive test for service providers will evolve in parallel. Models, engineering assets and production capabilities remain necessary,



Executive Summary

but possessing them will explain progressively less of the difference among stronger competitors as foundational capabilities become more widely available. Differentiation will increasingly come from codifying enterprise context and workflow knowledge, integrating AI into heterogeneous technology environments and carrying implementation learning from one engagement into the next.

A more fundamental challenge is emerging beyond AI production itself, as enterprises become increasingly capable of building, governing and operating GenAI while evidence of equivalent progress in redesigning roles, workflows and operating models remains more limited. Organizational absorption will therefore become increasingly important as machine interpretation, synthesis and knowledge access become cheaper and more widely available.

The defining question for the next phase of GenAI is no longer where machine cognition can be added to enterprise work, but which parts of enterprise work should continue to operate as though cognition were scarce.

For large and midsize providers, the competitive frontier is consequently shifting from delivering GenAI capability toward repeatedly converting increasingly abundant machine cognition into meaningful enterprise change.

GenAI is changing the economics of work faster than enterprises are redesigning work around those new economics.

Repetition is not repeatability. Deeper industrialization begins when providers reuse knowledge of enterprise context, workflows, integrations, evaluation and controls, reducing the need to reinvent successful GenAI implementations for each engagement.





Provider Positioning

Page 1 of 6

	Strategy and Consulting Services — Large	Strategy and Consulting Services — Midsize	Development and Deployment Services — Large	Development and Deployment Services — Midsize
Accenture	Leader	Not In	Leader	Not In
Accion Labs	Not In	Leader	Not In	Leader
Akkodis	Product Challenger	Not In	Product Challenger	Not In
Altimetrik	Not In	Product Challenger	Not In	Product Challenger
Apexon	Not In	Leader	Not In	Leader
Ascendion	Not In	Leader	Not In	Leader
Aspire Systems	Not In	Contender	Not In	Product Challenger
Atain	Not In	Contender	Not In	Contender
Atos	Rising Star ★	Not In	Leader	Not In
Birlasoft	Not In	Leader	Not In	Leader





Provider Positioning

Page 2 of 6

	Strategy and Consulting Services — Large	Strategy and Consulting Services — Midsize	Development and Deployment Services — Large	Development and Deployment Services — Midsize
Brillio	Not In	Leader	Not In	Leader
Capgemini	Leader	Not In	Leader	Not In
CGI	Contender	Not In	Contender	Not In
CI&T	Not In	Product Challenger	Not In	Product Challenger
Coforge	Not In	Product Challenger	Not In	Product Challenger
Cognizant	Leader	Not In	Leader	Not In
Cybage	Not In	Leader	Not In	Leader
Deloitte	Leader	Not In	Market Challenger	Not In
EPAM Systems	Contender	Not In	Contender	Not In
EXL	Not In	Leader	Not In	Leader





Provider Positioning

Page 3 of 6

	Strategy and Consulting Services — Large	Strategy and Consulting Services — Midsize	Development and Deployment Services — Large	Development and Deployment Services — Midsize
Firstsource	Not In	Leader	Not In	Leader
Fujitsu	Market Challenger	Not In	Market Challenger	Not In
Genpact	Product Challenger	Not In	Product Challenger	Not In
Happiest Minds	Not In	Product Challenger	Not In	Product Challenger
HCLTech	Leader	Not In	Leader	Not In
Hexaware	Not In	Leader	Not In	Leader
HTC Global Services	Not In	Rising Star ★	Not In	Rising Star ★
IBM	Leader	Not In	Leader	Not In
Indium Software	Not In	Product Challenger	Not In	Product Challenger
Infosys	Leader	Not In	Leader	Not In





Provider Positioning

Page 4 of 6

	Strategy and Consulting Services — Large	Strategy and Consulting Services — Midsize	Development and Deployment Services — Large	Development and Deployment Services — Midsize
Innova Solutions	Not In	Product Challenger	Not In	Product Challenger
ITC Infotech	Not In	Product Challenger	Not In	Product Challenger
Konecta	Not In	Leader	Not In	Leader
Kyndryl	Product Challenger	Not In	Contender	Not In
LTM	Product Challenger	Not In	Product Challenger	Not In
Mastek	Not In	Contender	Not In	Contender
MHP	Not In	Market Challenger	Not In	Market Challenger
Mindsprint	Not In	Contender	Not In	Contender
Movate	Not In	Product Challenger	Not In	Product Challenger
NTT DATA	Leader	Not In	Leader	Not In





Provider Positioning

Page 5 of 6

	Strategy and Consulting Services — Large	Strategy and Consulting Services — Midsize	Development and Deployment Services — Large	Development and Deployment Services — Midsize
Orange Business	Product Challenger	Not In	Product Challenger	Not In
Orion Innovation	Not In	Product Challenger	Not In	Product Challenger
Persistent Systems	Not In	Leader	Not In	Leader
Qburst	Not In	Contender	Not In	Contender
Rackspace Technology	Product Challenger	Not In	Product Challenger	Not In
Softtek	Not In	Leader	Not In	Leader
Sopra Steria	Contender	Not In	Contender	Not In
Stefanini	Not In	Product Challenger	Not In	Product Challenger
Syntax	Not In	Product Challenger	Not In	Product Challenger
TCS	Leader	Not In	Leader	Not In





Provider Positioning

Page 6 of 6

	Strategy and Consulting Services — Large	Strategy and Consulting Services — Midsize	Development and Deployment Services — Large	Development and Deployment Services — Midsize
Tech Mahindra	Product Challenger	Not In	Rising Star ★	Not In
TP	Market Challenger	Not In	Market Challenger	Not In
Trigent	Not In	Leader	Not In	Leader
T-Systems	Product Challenger	Not In	Product Challenger	Not In
Unisys	Not In	Leader	Not In	Leader
Virtusa	Not In	Leader	Not In	Leader
Wipro	Leader	Not In	Leader	Not In
Xoriant	Not In	Contender	Not In	Contender
YASH Technologies	Not In	Product Challenger	Not In	Product Challenger
Zensar Technologies	Not In	Market Challenger	Not In	Market Challenger



The study provides insights into evolving market trends and competitive dynamics among providers of **GenAI services**.

Simplified Illustration Source: ISG 2026

Strategy and Consulting Services — Large

Strategy and Consulting Services — Midsize

Development and Deployment Services — Large

Development and Deployment Services — Midsize

Definition

GenAI is redefining the enterprise technology landscape in 2026, evolving from experimentation to scaled, value-driven transformation. It now serves as a strategic enabler of intelligent enterprises, reshaping operations and accelerating innovation cycles.

Today, GenAI powers not only text, code, images and video but also autonomous workflows, decision intelligence and hyperpersonalized customer experiences across industries.

Key trends shaping GenAI investments include:

- As organizations move beyond pilots, advisory services play a critical role in defining GenAI roadmaps, establishing governance models and building scalable operating frameworks to accelerate ROI and enterprise-wide adoption.
- LLMs are advancing in reasoning, contextual awareness and multimodal capabilities, while smaller, domain-specific models gain traction for cost-efficient, high-precision applications. Fine-tuning and retrieval-augmented generation (RAG) enable secure, data-driven customization.

- As deployments scale, LLMOps frameworks manage model lifecycle, versioning and security, ensuring traceability. Continuous evaluation, prompt refinement and bias detection help sustain accuracy and business alignment.
- Autonomous AI agents are increasingly capable of executing multi-step tasks, orchestrating workflows across enterprise applications and augmenting decision-making.
- Enterprises are embedding transparency, explainability and risk management into GenAI strategies to safeguard trust and ensure sustainable adoption.

While GenAI's potential is transformative, success depends on scalable architecture, disciplined governance and continuous optimization to unlock long-term value.



Scope of the Report

This ISG Provider Lens® quadrant report covers the following four quadrants for services: Strategy and Consulting Services — Large, Strategy and Consulting Services — Midsize, Development and Deployment Services — Large and Development and Deployment Services — Midsize.

This ISG Provider Lens® study offers IT decision-makers:

- Transparency on the strengths and weaknesses of relevant providers
- A differentiated positioning of providers by segments (quadrants)
- Focus on the Global market

Our study serves as the basis for important decision-making by covering providers' positioning, key relationships and go-to-market considerations. ISG advisors and enterprise clients also use information from these reports to evaluate their existing vendor relationships and potential engagements.

Provider Classifications

The provider position reflects the suitability of providers for a defined market segment (quadrant). Without further additions, the position always applies to all company sizes classes and industries. In case the service requirements from enterprise customers differ and the spectrum of providers operating in the local market is sufficiently wide, a further differentiation of the providers by performance is made according to the target group for products and services. In doing so, ISG either considers the industry requirements or the number of employees, as well as the corporate structures of customers and positions providers according to their focus area. As a result, ISG differentiates them, if necessary, into two client target groups that are defined as follows:

- **Large Providers:** Are those with revenues exceeding \$4 billion and a workforce of over 100,000 employees. They cater to multiple verticals, often spreading their resources across a broad range of industries. Their primary focus lies in serving large enterprises, often engaging in large transformation projects that require deep expertise, extensive resources, and the ability to manage complex, enterprise-wide

innovations. Their deep industry experience, broad service capabilities, and strategic partnerships with technology giants position them as key players in the global digital services landscape.

- **Midsize Providers:** On the other hand, generate less than \$4 billion in revenue and typically specialize in 3-4 verticals where they hold strong capabilities and significant revenue share. With a leaner workforce of under 75,000 employees, these providers adopt an agile and flexible approach, making them well-suited to serve both large enterprises and mid-market clients with tailored, industry-specific solutions. They also have strong inherent capabilities and heritage in Digital Engineering services. This combination of domain expertise, flexibility, and a strong focus on innovation positions them as effective partners for businesses seeking to implement cuttingedge technologies with a faster, more agile approach.
- **Specialists:** Are service providers uniquely positioned due to their niche capabilities, which are either deeply embedded in specific verticals (e.g., healthcare, financial services) or concentrated on specialized service areas like

AI and analytics. Typically, these providers focus intensely on 2-3 verticals where they hold a significant market share and expertise, allowing them to deliver highly tailored and innovative solutions. With a workforce of fewer than 10,000 employees, specialists leverage their agility and flexibility to serve both large and mid-market enterprises. Their approach emphasizes solution-based problem-solving, making them highly responsive to the specific needs of their clients.

The ISG Provider Lens® quadrants are created using an evaluation matrix containing four segments (Leader, Product & Market Challenger and Contender), and the providers are positioned accordingly. Each ISG Provider Lens® quadrant may include a service provider(s) which ISG believes has strong potential to move into the Leader quadrant. This type of provider can be classified as a Rising Star.

- **Number of providers in each quadrant:** ISG rates and positions the most relevant providers according to the scope of the report for each quadrant and limits the maximum of providers per quadrant to 25 (exceptions are possible).





Provider Classifications: Quadrant Key

Product Challengers offer a product and service portfolio that reflect excellent service and technology stacks. These providers and vendors deliver an unmatched broad and deep range of capabilities. They show evidence of investing to enhance their market presence and competitive strengths.

Contenders offer services and products meeting the evaluation criteria that qualifies them to be included in the IPL quadrant. These promising service providers or vendors show evidence of rapidly investing in products/ services and follow sensible market approach with a goal of becoming a Product or Market Challenger within 12 to 18 months.

Leaders have a comprehensive product and service offering, a strong market presence and established competitive position. The product portfolios and competitive strategies of Leaders are strongly positioned to win business in the markets covered by the study. The Leaders also represent innovative strength and competitive stability.

Market Challengers have a strong presence in the market and offer a significant edge over other vendors and providers based on competitive strength. Often, Market Challengers are the established and well-known vendors in the regions or vertical markets covered in the study.

★ **Rising Stars** have promising portfolios or the market experience to become a Leader, including the required roadmap and adequate focus on key market trends and customer requirements. Rising Stars also have excellent management and understanding of the local market in the studied region. These vendors and service providers give evidence of significant progress toward their goals in the last 12 months. ISG expects Rising Stars to reach the Leader quadrant within the next 12 to 24 months if they continue their delivery of above-average market impact and strength of innovation.

Not in means the service provider or vendor was not included in this quadrant. Among the possible reasons for this designation: ISG could not obtain enough information to position the company; the company does not provide the relevant service or solution as defined for each quadrant of a study; or the company did not meet the eligibility criteria for the study quadrant. Omission from the quadrant does not imply that the service provider or vendor does not offer or plan to offer this service or solution.





Strategy and Consulting Services — Large

Who Should Read This Section

This report is valuable for large service providers offering generative AI (GenAI) strategy and consulting services globally to understand their market position and for enterprises looking to evaluate these providers. The report highlights how large service providers enable enterprises to scale GenAI from pilots to production through operating models, governance frameworks and AI maturity road maps. It highlights how providers advance use case expansion and strengthen data and risk governance to deliver measurable business value.

Chief information and compliance officers

Chief information and compliance officers should read this report to identify providers that embed responsible AI frameworks, model explainability and audit trails into enterprise-wide GenAI strategies. The report highlights providers that can align governance with evolving regulations and mitigate risks of data leakage, model inversion and unauthorized access across hybrid and sovereign deployments.

Chief data and AI officers

Chief data and AI officers should read this report to identify providers that design enterprise-scale data governance and evaluation strategies for GenAI. The report highlights how providers enable robust AI governance, operationalize GenAI across the enterprise, optimize model performance and cost, and support secure, flexible deployment options.

Line-of-business managers

Line-of-business managers should read this report to understand how providers translate GenAI strategy into measurable business outcomes such as productivity, personalization and customer engagement. The report offers insight into how providers help business leaders build repeatable, ROI-driven roadmaps that move initiatives beyond pilots into scaled, orchestrated enterprise deployments.


Strategy professionals

Strategy professionals should read this report to understand how providers help enterprises translate GenAI investments into long-term competitive advantage. The report highlights how providers combine domain expertise, GenAI advisory and orchestration frameworks to build scalable, resilient organizations aligned with evolving market and regulatory dynamics.



Generative AI Services Strategy and Consulting Services – Large

Global 2026



This quadrant assesses large providers guiding enterprise-wide **GenAI transformation** through strategy, portfolio prioritization, **governance, architecture**, operating-model and **workflow redesign** to deliver **measurable value**.

Sushree Acharya



Strategy and Consulting Services – Large

Definition

In this quadrant, ISG evaluates providers delivering strategy and consulting services to help enterprises plan, govern and operationalize GenAI initiatives. These services equip business and IT leaders with decision frameworks and practical guidance to set priorities, align stakeholders and translate GenAI ambition into executable roadmaps.

Providers assess use case identification and prioritization, data readiness and risk and advise on model selection, fine-tuning strategies and scalable architecture design. They define governance frameworks, responsible AI policies and compliance structures to manage regulatory and ethical risks. Increasingly, consultants define LLMOps-enabled operating models, lifecycle management processes and continuous evaluation mechanisms to sustain performance at scale.

Providers integrate strategy, governance and organizational readiness with change enablement approaches to drive industrialized GenAI adoption and sustained value realization.

Eligibility Criteria

1. Demonstrate **domain-specific advisory experience** in designing and scaling enterprise-wide GenAI strategies
2. Provide client references showcasing **PoC-to-production progression, use case prioritization** and value creation, with a value-tracking approach
3. Exhibit advisory expertise in **LLMs, small language models (SLMs), fine-tuning, RAG architectures and model selection strategies**
4. Demonstrate advisory capabilities in defining **LLMOps frameworks**, including lifecycle management, monitoring, versioning and governance
5. Showcase advisory capabilities in framing **continuous evaluation mechanisms**, including performance monitoring, bias detection and model optimization frameworks
6. Present evidence of transitioning from **strategy to implementation** readiness, beyond PoCs
7. Demonstrate partnerships across **technology providers, cloud hyperscalers, academia and startup ecosystems**
8. Showcase investments in reusable advisory-related **proprietary assets such as accelerators, frameworks and platforms**
9. Establish and implement robust **responsible AI and governance frameworks**, addressing transparency, explainability, security and regulatory compliance
10. Support **operating model design, change management and AI capability building** to enable sustainable GenAI adoption



Strategy and Consulting Services – Large

Observations

Large IT service providers are shifting GenAI consulting from use case identification and experimentation toward enterprise-wide business and operating model transformation. The demand focuses on prioritizing AI initiatives by business value, feasibility and scalability, alongside measurable outcomes and end-to-end business process redesign. Providers are responding with structured transformation frameworks, maturity assessments, value realization methodologies and roadmaps connecting strategy with engineering, integration and operations, while favoring architectures that preserve existing business rules and systems of record rather than displace them. Governance is moving from policy-level advisory into operating requirements, embedded as risk controls, security and compliance within programs, alongside mandates around data sovereignty, model control and multicloud environments. Economics is becoming a design input, with token consumption and ROI folded into business case development. Commercial

models are shifting away from headcount toward outcome-, transaction- and consumption-based pricing. Providers are codifying consulting expertise into reusable, AI-enabled IP and industry-specific assets to reduce dependence on bespoke advisory. Ecosystem and foundation model partnerships mature from technology access into joint engineering and commercialization, supported by structured change management and workforce enablement. Large providers are positioning GenAI consulting as a continuous transformation capability, fusing strategy, governance, economics and execution for scaled, outcome-led adoption.

From the 111 companies assessed for this study, 24 qualified for this quadrant, with 10 being Leaders and One Rising Star.

accenture

Accenture brings a strong enterprise re-invention lens to GenAI strategy, connecting technology ambition with operating model change, responsible AI, workforce readiness and measurable business outcomes across a broad industry footprint and transformation agenda.

Capgemini

Capgemini's RAISE 2.0 and human-AI chemistry underpin an enterprise GenAI strategy that combines strategic advisory, operating model redesign, engineering and governance to drive AI transformation at scale.

cognizant

Cognizant's AI Builder and BASIS frameworks provide a structured consulting approach for translating enterprise priorities into redesigned operating models, governed solutions and scalable AI implementation.

Deloitte.

Deloitte's GenAI Navigator and IndustryAdvantage™ translate GenAI ambitions into financially linked roadmaps, operating models and value cases, supported by governance, workforce redesign and Trustworthy AI™ advisory for disciplined GenAI adoption.

HCLTech

HCLTech has productized its AI advisory lifecycle, linking readiness, governance architecture and operating model services with engineering pathways that move GenAI programs from strategic alignment toward controlled implementation and scale.

IBM

IBM's Enterprise Advantage combines consulting expertise with reusable AI assets to help enterprises scale AI-led transformation, embedding strategy, technology and execution into a more repeatable, software-led delivery model.



Strategy and Consulting Services — Large



Infosys is integrating Topaz Fabric, AI Value Canvas and responsible AI assets into a consulting model that links value discovery, process transformation, governance and enterprise architecture to drive structured, business-led AI adoption.



NTT DATA's global AI unit and AIVista investment unite strategy, product engineering and commercialization, strengthening AI-native innovation and accelerating industry-led, outcome-driven GenAI adoption.



TCS strengthened its AI consulting capabilities through a dedicated transformation unit that combines its Human + AI autonomy service model with strategic hyperscaler and frontier model partnerships.



Wipro's expansion of WEGA and WINGS reinforces a platform-led transformation model that connects AI strategy, governance and process redesign with downstream execution, creating continuity between advisory engagements and GenAI program operationalization.



Atos (Rising Star) has expanded Polaris AI into a reusable enterprise foundation to strengthen its GenAI strategy and consulting capabilities, spanning AI governance, modernization, engineering, security and controlled deployment across regulated environments.





Strategy and Consulting Services — Midsize

Who Should Read This Section

This report is valuable for midsize service providers offering generative AI (GenAI) strategy and consulting services globally to understand their market position and for enterprises looking to evaluate these providers. The report highlights how midsize service providers differentiate through specialized AI capabilities, domain expertise and agile delivery. It highlights how they help enterprises scale GenAI use cases, strengthen data and AI governance, integrate AI into workflows and accelerate business value.

Chief information and compliance officers

Chief information and compliance officers should read this report to identify providers that embed responsible AI frameworks, model explainability and audit trails into enterprise-wide GenAI strategies. It highlights providers that can align governance with evolving regulations and mitigate risks of data leakage, model inversion and unauthorized access across hybrid and sovereign deployments.

Chief data and AI officers

Chief data and AI officers should read this report to identify providers that design enterprise-scale data governance and evaluation strategies for GenAI. The report highlights how providers enable robust AI governance, operationalize GenAI across the enterprise, optimize model performance and cost, and support secure, flexible deployment options.


Line-of-business managers

Line-of-business managers should read this report to understand how providers translate GenAI strategy into measurable business outcomes such as productivity, personalization and customer engagement. The report offers insight into how providers help business leaders build repeatable, ROI-driven roadmaps that move initiatives beyond pilots into scaled, orchestrated enterprise deployments.

Strategy professionals

Strategy professionals should read this report to understand how providers help enterprises translate GenAI investments into long-term competitive advantage. The report highlights how providers combine domain expertise, GenAI advisory and orchestration frameworks to build scalable, resilient organizations aligned with evolving market and regulatory dynamics.





This quadrant assesses providers guiding enterprises from GenAI adoption to measurable value through strategy, portfolio prioritization, governance, architecture, enterprise integration and workflow transformation.

Gowtham Kumar Sampath



Strategy and Consulting Services — Midsize

Definition

In this quadrant, ISG evaluates providers delivering strategy and consulting services to help enterprises plan, govern and operationalize GenAI initiatives. These services equip business and IT leaders with decision frameworks and practical guidance to set priorities, align stakeholders and translate GenAI ambition into executable roadmaps.

Providers assess use case identification and prioritization, data readiness and risk and advise on model selection, fine-tuning strategies and scalable architecture design. They define governance frameworks, responsible AI policies and compliance structures to manage regulatory and ethical risks. Increasingly, consultants define LLMOps-enabled operating models, lifecycle management processes and continuous evaluation mechanisms to sustain performance at scale.

Providers integrate strategy, governance and organizational readiness with change enablement approaches to drive industrialized GenAI adoption and sustained value realization.

Eligibility Criteria

1. Demonstrate **domain-specific advisory experience** in designing and scaling enterprise-wide GenAI strategies
2. Provide client references showcasing **PoC-to-production progression, use case prioritization** and value creation, with a value-tracking approach
3. Exhibit advisory expertise in **LLMs, small language models (SLMs), fine-tuning, RAG architectures and model selection strategies**
4. Demonstrate advisory capabilities in defining **LLMOps frameworks**, including lifecycle management, monitoring, versioning and governance
5. Showcase advisory capabilities in framing **continuous evaluation mechanisms**, including performance monitoring, bias detection and model optimization frameworks
6. Present evidence of transitioning from **strategy to implementation** readiness, beyond PoCs
7. Demonstrate partnerships across **technology providers, cloud hyperscalers, academia and startup ecosystems**
8. Showcase investments in reusable advisory-related **proprietary assets such as accelerators, frameworks and platforms**
9. Establish and implement robust **responsible AI and governance frameworks**, addressing transparency, explainability, security and regulatory compliance
10. Support **operating model design, change management and AI capability building** to enable sustainable GenAI adoption



Strategy and Consulting Services – Midsize

Observations

Midsize providers are strengthening GenAI consulting around focused transformation priorities, combining strategy and use-case selection with implementation knowledge, domain expertise and organizational readiness. Their propositions are increasingly moving beyond experimentation toward identifying where GenAI can produce measurable value, establish governance and help enterprises progress from individual initiatives to repeatable adoption.

A key opportunity lies in more closely connecting GenAI strategy to work design. Enterprises are embedding GenAI into knowledge-intensive activities considerably faster than they are redesigning workflows, roles and handoffs around increasingly accessible machine cognition. Midsize providers can differentiate through focused process and domain expertise that helps clients determine where augmentation is sufficient and where the economics justify deeper workflow change.

Governance, data and knowledge readiness, operating models and change enablement are therefore becoming integral to consulting engagements. Midsize providers are also using reusable frameworks, accelerators and specialized expertise to reduce strategy-to-implementation gaps without requiring large-scale transformation programs. Stronger propositions combine this implementation pragmatism with measurable business objectives, helping enterprises prioritize investments, establish scalable controls and translate GenAI opportunities into organizational changes that can be realistically absorbed.

From the 111 companies assessed for this study, 36 qualified for this quadrant, with 15 being Leaders and one Rising Star.

Accion Labs

Accion Labs has industrialized its AI engineering portfolio around SPEX, ECL, KnowledgeMesh and ASIMOV, supported by production deployments across knowledge management, automation and modernization.



Apexon has expanded its engineering-led AI portfolio through reusable domain assets, flexible model architectures and AgentRise controls that support governed production deployment.

Ascendion

Ascendion's AAVA 2.0 introduces governance and execution layers, reusable golden agents and developer tool integration, strengthening its engineering-led approach to production-scale enterprise AI delivery.



Birlasoft has expanded its engineering-led AI proposition through Cogito and reusable delivery assets, industrializing application development, orchestration, benchmarking and lifecycle controls across enterprise environments.

Brillio

Brillio has repositioned ADAM as an enterprise AI accelerator, strengthening an engineering-led delivery model built around reusable assets, governed deployment and AI-native talent.



Cybage has institutionalized ARGO and SmartPal as proprietary engineering foundations, combining reusable orchestration, retrieval, governance and multicloud deployment capabilities across its production-grade enterprise AI portfolio.



EXL is consolidating reusable engineering assets within EXLerate and its agentic factory, bringing together domain skills, guardrails, contextual data, model economics and portable AI components for production-oriented enterprise AI engineering.



Strategy and Consulting Services – Midsize



Firstsource has organized its engineering portfolio around Kairos and Agentic AI Studio, combining reusable domain assets, evaluation, orchestration and monitoring with a delivery model that extends from build to governed, managed AI operations at scale.

HEXAWARE

Hexaware has expanded its AI engineering portfolio around AgentVerse while retaining RapidX, Tensai and Amaze as anchors for software development, operations and modernization, creating a broader lifecycle spanning build, test, deployment, governance and run.

Konecta

Konecta has established Kolibri as the engineering anchor of its AI portfolio, combining an open orchestration layer with model abstraction, runtime governance, observability and FinOps as it industrializes production delivery across its global and regional teams.



Persistent Systems advanced SASVA into a team-oriented AI engineering harness and continued investment in iAURA and GenAI Hub, giving its enterprise AI development portfolio proprietary assets spanning software, data and governed AI operations at a global scale.

Softtek

Softtek has concentrated recent Enterprise AI investment in FRIDA, SALMA and ASE, pairing proprietary engineering assets with a production-first lifecycle and a growing emphasis on reusable industry patterns, AI-ready data and managed AI operations at scale.

Trigent

Trigent has made ArkOS the center of its enterprise AI engineering portfolio, combining integration, orchestration, governance and observability with AI-DLC as it moves from rapid prototypes to repeatable, production-oriented enterprise delivery.



Unisys is consolidating AI engineering around the Unisys Intelligence Accelerator and Service Experience Accelerator, creating reusable foundations for model integration, knowledge services, application engineering and intelligent managed operations.

virtusa

Virtusa's shift from AI-assisted delivery toward AI-native engineering has formalized a delivery model built around smaller human teams, specialized AI agents, lifecycle scaffolding and reusable Helio assets across development, testing, modernization and operations.



HTC Global Services (Rising Star) has consolidated reusable GenAI engineering components within MAiGE, combining data, prompt, retrieval-augmented generation (RAG), model and evaluation studios with deployment practices that move its AI portfolio toward more repeatable production engineering at scale.



Unisys



“Unisys centers its AI strategy around running enterprise technology estates, bringing operational context into architecture, governance and modernization decisions to shape AI programs from the outset for production use and sustained operations.”

Gowtham Kumar Sampath

Overview

Unisys is headquartered in Pennsylvania, U.S. It has more than 15,000 employees across 20 countries. In FY25, the company generated \$2.0 billion in revenue. Unisys is a global technology services provider applying enterprise AI across cloud, applications, digital workplace and managed operations. Its AI framework centers on the Develop–Transform–Orchestrate model, supported by AI and data readiness assessments, modernization disciplines and governance practices. Unisys combines proprietary assets, such as the Intelligence Accelerator, along with hardware, data and other ecosystem technologies, using its managed services footprint as the foundation for AI adoption.

Strengths

Readiness before deployment: Unisys assesses factors such as data readiness, AI operational readiness, application portfolios and development maturity. Teams evaluate foundations, technical debt, governance, security and human factors before setting priorities. The portfolio provides enterprises with structured entry points and keeps AI planning aligned with conditions that are needed to move initiatives beyond experimentation at scale.

Operations-grounded AI strategy: Unisys anchors AI advisory in the cloud, application, workplace and service estates that its teams already manage. Consulting connects readiness, architecture and modernization decisions with the operational conditions required for production use.

Unisys’s operations-grounded consulting approach emphasizes practical adoption to keep strategy linked with executable technology decisions.

Managed services continuity: Unisys connects consulting with implementation and ongoing operations across infrastructure, applications and digital workplace services. Its delivery model extends AI planning into modernization and managed execution without requiring a separate operating model. Responsibility for enterprise technology estates also provides an operational context to guide priorities, integration choices and subsequent service optimization.

Caution

Unisys’ strategy services remain anchored in technology operations, applications and managed services, with limited depth in business function redesign and industry-specific AI operating models. Also, its consulting position is the strongest when strategy is expected to progress into technology execution.





Generate Content

Development and Deployment Services — Large

Who Should Read This Section

This report is valuable for large service providers offering generative AI (GenAI) development and deployment services globally to understand their market position and for enterprises looking to evaluate these providers. The report highlights how large providers fine-tune LLMs, develop SLMs and build orchestration frameworks to industrialize GenAI, enabling secure, multimodal deployment at production scale across complex enterprise environments.

Line-of-business managers

Line-of-business managers should read this report to understand how providers build and deploy agentic platforms aligned with business goals. The report highlights providers' capabilities in orchestrating multimodal GenAI platforms, managing agent lifecycles and providing observability tooling. It also shows how these solutions improve efficiency, decision-making and customer engagement at scale.

Chief data and AI officers

Chief data and AI officers should read this report to identify providers that operationalize data governance and evaluation pipelines for GenAI. The report outlines how providers ensure data quality, security and compliance across the AI lifecycle. It also shows how hybrid LLM-SLM architectures and control towers deliver transparent, scalable oversight across GenAI applications, including private and sovereign deployments.

Chief information officers

Chief information officers should read this report to identify providers that ensure seamless GenAI deployments, with a focus on system interoperability and scalability across cloud and legacy environments. The report highlights how providers architect secure, sovereign-ready infrastructure that maintains data accuracy, performance and regulatory alignment at production scale.


Strategy professionals

Strategy professionals should read this report to understand how providers help enterprises translate GenAI investments into long-term competitive advantage. The report highlights how providers combine domain expertise, engineering depth and orchestration capabilities to build future-ready platforms that are scalable, resilient and responsive to market dynamics.



Generative AI Services Development and Deployment Services – Large

Global 2026



This quadrant assesses large providers engineering and operating production GenAI across complex enterprise estates, integrating context, applications and workflows with scalable architectures, continuous assurance and governance.

Sushree Acharya



Development and Deployment Services – Large

Definition

In this quadrant, ISG evaluates providers delivering development and deployment services to industrialize and scale GenAI solutions from pilot to production. Providers support the end-to-end lifecycle, including solution design, model customization, systems integration, deployment and ongoing optimization.

Providers architect secure, cost-efficient cloud and hybrid infrastructures to enable scalable model training and inference. They deliver fine-tuning foundation models, RAG architectures, multimodal integration and embedding GenAI into business workflows. Providers operationalize LLMOps practices, including automated pipelines, monitoring, version control, performance management and continuous evaluation to address drift and reliability at scale. Security controls, data governance and regulatory compliance mechanisms are increasingly core to delivery.

Providers deliver resilient, production-grade GenAI environments that accelerate value while ensuring stability, security and trust.

Eligibility Criteria

1. Demonstrate expertise in **designing and deploying enterprise-grade cloud and hybrid infrastructure** for scalable model training and inference
2. Optimize compute utilization through **cost-efficient orchestration, GPU management and performance tuning**
3. Showcase capabilities in building and **fine-tuning foundation models, implementing RAG architectures and developing and customizing SLMs** using domain-specific data
4. Maintain a skilled **GenAI engineering and data science team** with expertise in preprocessing, model optimization and evaluation
5. Operationalize **LLMOps pipelines**, including automated deployment, version control, monitoring and lifecycle management
6. Implement **continuous evaluation and performance monitoring** to address drift, hallucinations and reliability at scale
7. Demonstrate proficiency in **API integration, orchestration and embedding GenAI into enterprise applications and workflows**
8. Establish **security, data governance and regulatory compliance controls**, including encryption and access management
9. Showcase **production-grade use cases and scaled deployments beyond PoCs**, aligned to measurable ROI and business outcomes
10. Demonstrate capability in developing and deploying **multimodal, GenAI and agentic AI** solutions across enterprise environments



Development and Deployment Services – Large

Observations

Large IT service providers report GenAI development maturing from pilot-led experimentation toward industrialized engineering and production deployment, with demand concentrated in document-intensive, regulated and knowledge-intensive processes, software engineering, modernization and operational use cases. Providers are building reusable architectures, engineering blueprints, accelerators and factory-based delivery models to improve deployment speed, consistency and scale, replacing project-specific builds with componentized engineering.

Production environments are increasingly hybrid, pairing LLMs with traditional AI, deterministic business rules, workflow automation and enterprise systems. Retrieval augmented generation (RAG) is evolving from standalone retrieval to embedded contextual capability, underpinned by stronger data grounding and a growing discipline of context engineering, while model selection is becoming

model-agnostic, mixing foundation, small and domain-tuned models based on accuracy, cost, workload and sovereignty needs.

Operationalizing GenAI at scale is creating new engineering requirements around observability, security, governance and AI FinOps, with monitoring extending beyond model accuracy to inference cost, drift and workflow behavior. GenAI is being embedded across the software development lifecycle, supporting requirements interpretation, coding, testing, debugging and modernization. This is driving providers toward standardized engineering practices that balance performance, cost, control and scalability across enterprise environments.

From the 111 companies assessed for this study, 24 qualified for this quadrant, with 10 being Leaders and one Rising Star.

accenture

Accenture demonstrates strong GenAI industrialization capabilities through AI Refinery™ and Distiller, combining model flexibility, reusable engineering assets, structured deployment programs and managed evaluation to support enterprise-scale production adoption with repeatable delivery practices.

Atos

Atos extends Polaris AI suite across software delivery, modernization and testing, integrating AI with core delivery flows and reinforcing an engineering-led model focused on repeatable deployment, production control and enterprise-scale execution.

Capgemini

Capgemini's engineering-led AI implementation combines agent development, integration, governance and operations into a common foundation, supporting an industrialized approach to developing enterprise AI across heterogeneous cloud, model and technology environments.

cognizant

Cognizant industrializes enterprise-wide AI deployment through a combination of reusable development assets, production operations and platform-neutral integration with workforce initiatives designed to embed AI-native practices across delivery.



Development and Deployment Services – Large

HCLTech

HCLTech combines its execution assets with hyperscaler, model and enterprise-platform ecosystems, enabling clients with multiple routes to operationalize agents, preserving architecture choice while strengthening focus on controls and production integration.



IBM's integrated engineering flow connects process decomposition, product design and GenAI application development and reduces handoffs between discovery and implementation, preserving enterprise context throughout the build.



Infosys' core GenAI capabilities expand across contextual reasoning, enterprise knowledge, code understanding and document intelligence, translating generative models into production applications within enterprise workflows.



NTT DATA's full-stack AI engineering expertise across data, cloud, infrastructure and private environments, combined with its ability to orchestrate LLMs, AI platforms and agentic frameworks, enables enterprises to operationalize business-aligned GenAI solutions.



TCS' capabilities across model adaptation, contextual intelligence and multimodal AI, combined with foundation models and deep domain expertise, enable enterprises to develop industry-specific AI applications that enhance decision intelligence and drive business outcomes.



Wipro is strengthening its GenAI engineering foundation through WEGA and WINGS, creating a scalable platform backbone for build and run environments that support AI-driven software development, modernization and enterprise operations across delivery landscapes.



Tech Mahindra (Rising Star) has a robust GenAI engineering ecosystem through Orion, Swifter.IO, reusable agent frameworks and domain-specific models, enabling scalable AI development, governance, workflow integration and application modernization across enterprise delivery.





Generate Content

Development and Deployment Services — Midsize

Who Should Read This Section

This report is valuable for midsize service providers offering generative AI (GenAI) development and deployment services globally to understand their market position and for enterprises looking to evaluate these providers. The report highlights how midsize providers deliver modular, GenAI solutions tailored to industry workflows, enabling incremental AI adoption without extensive architectural transformation.

Line-of-business managers

Line-of-business managers should read this report to understand how providers build and deploy agentic platforms aligned with business goals. The report highlights providers' capabilities in orchestrating multimodal GenAI platforms, managing agent lifecycles and providing observability tooling. It also shows how these solutions improve efficiency, decision-making and customer engagement at scale.

Chief data and AI officers

Chief data and AI officers should read this report to identify providers that operationalize data governance and evaluation pipelines for GenAI. The report outlines how providers ensure data quality, security and compliance across the AI lifecycle. It also shows how hybrid LLM-SLM architectures and control towers deliver transparent, scalable oversight across GenAI applications, including private and sovereign deployments.

Chief information officers

Chief information officers should read this report to identify providers that ensure seamless GenAI deployment, with a focus on system interoperability and scalability across cloud and legacy environments. The report highlights how providers architect secure, sovereign-ready infrastructure that maintains data accuracy, performance and regulatory alignment at production scale.


Strategy professionals

Strategy professionals should read this report to understand how providers help enterprises translate GenAI investments into long-term competitive advantage. The report highlights how providers combine domain expertise, engineering depth and orchestration capabilities to build future-ready platforms that are scalable, resilient and responsive to market dynamics.



Generative AI Services Development and Deployment Services – Midsize

Global 2026



This quadrant assesses providers developing and deploying production GenAI, **integrating enterprise context**, applications and **workflows** with scalable architecture, **continuous assurance**, **governance** and measurable **outcomes**.

Gowtham Kumar Sampath



Development and Deployment Services – Midsize

Definition

In this quadrant, ISG evaluates providers delivering development and deployment services to industrialize and scale GenAI solutions from pilot to production. Providers support the end-to-end lifecycle, including solution design, model customization, systems integration, deployment and ongoing optimization.

Providers architect secure, cost-efficient cloud and hybrid infrastructures to enable scalable model training and inference. They deliver fine-tuning foundation models, RAG architectures, multimodal integration and embedding GenAI into business workflows. Providers operationalize LLMOps practices, including automated pipelines, monitoring, version control, performance management and continuous evaluation to address drift and reliability at scale. Security controls, data governance and regulatory compliance mechanisms are increasingly core to delivery.

Providers deliver resilient, production-grade GenAI environments that accelerate value while ensuring stability, security and trust.

Eligibility Criteria

1. Demonstrate expertise in **designing and deploying enterprise-grade cloud and hybrid infrastructure** for scalable model training and inference
2. Optimize compute utilization through **cost-efficient orchestration, GPU management and performance tuning**
3. Showcase capabilities in building and **fine-tuning foundation models, implementing RAG architectures and developing and customizing SLMs** using domain-specific data
4. Maintain a skilled **GenAI engineering and data science team** with expertise in preprocessing, model optimization and evaluation
5. Operationalize **LLMOps pipelines**, including automated deployment, version control, monitoring and lifecycle management
6. Implement **continuous evaluation and performance monitoring** to address drift, hallucinations and reliability at scale
7. Demonstrate proficiency in **API integration, orchestration and embedding GenAI into enterprise applications and workflows**
8. Establish **security, data governance and regulatory compliance controls**, including encryption and access management
9. Showcase **production-grade use cases and scaled deployments beyond PoCs**, aligned to measurable ROI and business outcomes
10. Demonstrate capability in developing and deploying **multimodal, GenAI and agentic AI** solutions across enterprise environments



Development and Deployment Services – Midsize

Observations

Midsize providers are expanding GenAI development from standalone applications and generic RAG toward contextualized solutions embedded in enterprise workflows. Their strengths increasingly lie in focused engineering, domain specialization and reusable implementation assets that can accelerate deployment while adapting solutions to specific enterprise requirements.

Enterprise context is becoming an important source of differentiation as GenAI progresses from retrieval toward interpretation and workflow participation. Providers are combining enterprise knowledge, metadata, permissions and domain context with GenAI, conventional AI, deterministic technologies, APIs and applications rather than relying on models alone. This favors midsize providers that can combine technical depth with a close understanding of particular processes and technology environments.

Production capability is also expanding to include continuous evaluation, observability, governance, security and cost-performance

management. At the same time, accelerators, reusable code and AI factories increasingly establish delivery competence rather than differentiation on their own. Stronger midsize providers are moving toward contextual components, integrations, workflow patterns, evaluation and controls so that implementation learning carries across engagements. Their competitive opportunity lies in converting focused engineering and specialization into repeatable enterprise implementation rather than competing primarily on portfolio breadth or the volume of GenAI assets.

From the 111 companies assessed for this study, 36 qualified for this quadrant, with 15 being Leaders and one Rising Star.

Accion

Accion Labs has industrialized its AI engineering portfolio around SPEX, ECL, KnowledgeMesh and ASIMOV, supported by production deployments across knowledge management, automation and modernization.



Apexon has expanded its engineering-led AI portfolio through reusable domain assets, flexible model architectures and AgentRise controls that support governed production deployment.

Ascendion

Ascendion's AAVA 2.0 introduces governance and execution layers, reusable golden agents and developer tool integration, strengthening its engineering-led approach to production-scale enterprise AI delivery.



Birlasoft has expanded its engineering-led AI proposition through Cogito and reusable delivery assets, industrializing application development, orchestration, benchmarking and lifecycle controls across enterprise environments.

Brillio

Brillio has repositioned ADAM as an enterprise AI accelerator, strengthening an engineering-led delivery model built around reusable assets, governed deployment and AI-native talent.



Cybage has institutionalized ARGO and SmartPal as proprietary engineering foundations, combining reusable orchestration, retrieval, governance and multicloud deployment capabilities across its production-grade enterprise AI portfolio.



EXL is consolidating reusable engineering assets within EXLerate and its agentic factory, bringing together domain skills, guardrails, contextual data, model economics and portable AI components for production-oriented enterprise AI engineering.



Development and Deployment Services – Midsize



Firstsource has organized its engineering portfolio around Kairos and Agentic AI Studio, combining reusable domain assets, evaluation, orchestration and monitoring with a delivery model that extends from build to governed, managed AI operations at scale.

HEXAWARE

Hexaware has expanded its AI engineering portfolio around AgentVerse while retaining RapidX, Tensai and Amaze as anchors for software development, operations and modernization, creating a broader lifecycle spanning build, test, deployment, governance and run.

Konecta

Konecta has established Kolibri as the engineering anchor of its AI portfolio, combining an open orchestration layer with model abstraction, runtime governance, observability and FinOps as it industrializes production delivery across its global and regional teams.



Persistent Systems advanced SASVA into a team-oriented AI engineering harness and continued investment in iAURA and GenAI Hub, giving its enterprise AI development portfolio proprietary assets spanning software, data and governed AI operations at a global scale.

Softtek

Softtek has concentrated recent Enterprise AI investment in FRIDA, SALMA and ASE, pairing proprietary engineering assets with a production-first lifecycle and a growing emphasis on reusable industry patterns, AI-ready data and managed AI operations at scale.

Trigent

Trigent has made ArkOS the center of its enterprise AI engineering portfolio, combining integration, orchestration, governance and observability with AI-DLC as it moves from rapid prototypes to repeatable, production-oriented enterprise delivery.



Unisys is consolidating AI engineering around the Unisys Intelligence Accelerator and Service Experience Accelerator, creating reusable foundations for model integration, knowledge services, application engineering and intelligent managed operations.

virtusa

Virtusa's shift from AI-assisted delivery toward AI-native engineering has formalized a delivery model built around smaller human teams, specialized AI agents, lifecycle scaffolding and reusable Helio assets across development, testing, modernization and operations.



HTC Global Services (Rising Star) has consolidated reusable GenAI engineering components within MAiGE, combining data, prompt, retrieval-augmented generation (RAG), model and evaluation studios with deployment practices that move its AI portfolio toward more repeatable production engineering at scale.



Unisys



Leader

“Unisys operationalizes AI into the systems where enterprise work is executed, combining application engineering, service integration and proprietary operational platforms to move GenAI from isolated tools into secure production workflows.”

Gowtham Kumar Sampath

Overview

Unisys is headquartered in Pennsylvania, U.S. It has more than 15,000 employees across 20 countries. In FY25, the company generated \$2.0 billion in revenue. Unisys applies enterprise AI engineering across application modernization, digital workplace, knowledge services, cloud and intelligent operations. Its development portfolio combines reusable assets such as the Unisys Intelligence Accelerator (UIA) and Service Experience Accelerator (SEA) with hyperscaler, data and enterprise-platform technologies. Engineering teams emphasize integration with existing estates, secure deployment and managed-service continuity, giving the practice a production-oriented identity grounded in enterprise technology operations.

Strengths

Production-focused AI engineering:

Unisys embeds AI into applications, service management and knowledge workflows within established enterprise technology estates. Its teams connect models with source systems, developer tooling, ITSM platforms and operational data instead of isolating them behind standalone interfaces. Unisys' development capabilities focus on production integration, with enterprise systems and workflow execution playing a central role.

Knowledge workflow engineering: Unisys integrates retrieval, generation, clustering and knowledge curation into enterprise service and information workflows. Its engineering model addresses access to existing knowledge and creation or refinement of

operational content. Its integrations with service-management processes extend the capability beyond generic chat and establish knowledge engineering as a recurring layer across its workplace and managed services.

Reusable operational platforms: Unisys has developed shared engineering foundations for AI integration, enterprise context, knowledge services and orchestration. Accelerators such as UIA and SEA deliver reusable capabilities around models, data, APIs, guardrails and operational workflows while retaining links with client platforms. Engineering teams gain a common deployment foundation without relying on proprietary technology.

Caution

Unisys' development patterns remain concentrated in applications, IT operations, workplace services and enterprise knowledge. GenAI engineering across client-facing business processes and industry-specific functional workflows is less central to the portfolio.





Appendix

The ISG Provider Lens® 2026 – Generative AI Services study analyzes the relevant service providers in the global market, based on a multi-phased research and analysis process, and positions these providers based on the ISG Research methodology.

Study Sponsor:

Namratha Dharshan

Lead Authors:

Gowtham Kumar Sampath and
Sushree Acharya

Editors:

Ritu Sharma and Priyanka Richi

Data Analyst:

Rajesh Chillappagari

Consultant Advisors:

Loren Absher, Olga Kupriyanova, Ryan Hamze,
Dr. Dorotea Baljević, Diwahar Jawahar, and
Ritwik Dey

Project Manager:

Sukanya Nair

Information Services Group Inc. is solely responsible for the content of this report. Unless otherwise cited, all content, including illustrations, research, conclusions, assertions and positions contained in this report were developed by, and are the sole property of, Information Services Group Inc.

The research and analysis presented in this report includes research from the ISG Provider Lens® program, ongoing ISG Research programs, interviews with ISG advisors, briefings with service providers and analysis of publicly available market information from multiple sources. The data collected for this report represent information that ISG believes to be current as of September 2026 for providers that actively participated and for providers that did not. ISG recognizes that many mergers and acquisitions may have occurred since then, but this report does not reflect these changes.

All revenue references are in U.S. dollars (\$US) unless noted otherwise.

The study was conducted in the following steps:

1. Definition of Generative AI Services market
2. Use of questionnaire-based surveys of service providers/ vendor across all trend topics
3. Interactive discussions with service providers/vendors on capabilities and use cases
4. Leverage ISG's internal databases and advisor knowledge & experience (wherever applicable)
5. Detailed analysis and evaluation of services and service documentation based on the facts & figures received from providers and other sources.
6. Use of the following key evaluation criteria:
 - * Strategy and vision
 - * Innovation
 - * Brand awareness and presence in the market
 - * Sales and partner landscape
 - * Breadth and depth of portfolio of services offered
 - * Technology advancements



Author and Editor Biographies



Lead Author

Gowtham Kumar Sampath
Assistant Director and Principal Analyst

Gowtham Sampath is an Assistant Director with ISG Research, responsible for authoring ISG Provider Lens® quadrant reports for Banking Technology/Platforms, Digital Banking Services, Cybersecurity and Analytics Solutions & Services market.

With 15 years of market research experience, Gowtham works on analyzing and bridging the gap between data analytics providers and businesses, addressing market opportunities and best practices. In his role, he also works with advisors in addressing enterprise clients' requests for ad-hoc research requirements within the IT services sector, across industries.

He is also authoring thought leadership research, whitepapers, articles on emerging technologies within the banking sector in the areas of automation, DX and UX experience as well as the impact of data analytics across different industry verticals.



Lead Author

Sushree Acharya
Senior Analyst

Sushree Acharya is a Senior Analyst with ISG Research, responsible for authoring ISG Provider Lens® reports focused on Data & Analytics and AI. With nearly 9 years of experience across research, consulting and advisory roles, she specializes in market intelligence, hyperscaler and provider ecosystems, competitive positioning and go-to-market strategy within the technology services landscape.

Her work focuses on analyzing market evolution, identifying growth opportunities, competitive intelligence and evaluating emerging trends. She contributes to thought leadership, provider assessments, market reports and strategic insights that

help enterprises and service providers make informed business and investment decisions. Her experience also spans cloud and infrastructure services, as well as broader technology ecosystems.



Author and Editor Biographies




Study Sponsor

Namratha Darshan
Chief Business Leader

As a Chief Business Leader at ISG, Namratha Dharshan spearheads the BPO, AI and Analytics arm of the ISG Provider Lens® program, contributing to more than 20 reports. Under the aegis of this program, where she heads a team of analysts, Namratha manages the delivery of research findings on service provider intelligence. As a part of her role in the Senior Leadership Council, Namratha is the designated representative of the ISG India Research team, comprising more than 100 dynamic research professionals. In addition, Namratha is a speaker in ISG's flagship quarterly call, ISG Index™.

As a principal industry analyst and thought leader, Namratha is well recognized for her contributions to service provider intelligence and her understanding of the customer experience landscape, particularly the area of contact center services. She has also authored reports on other horizontal service lines such as finance and accounting and penned vertical focused reports for insurance.



IPL Product Owner

Jan Erik Aase
Partner and Global Head – ISG Provider Lens®

Mr. Aase brings extensive experience in the implementation and research of service integration and management of both IT and business processes. With over 35 years of experience, he is highly skilled at analyzing vendor governance trends and methodologies, identifying inefficiencies in current processes, and advising the industry. Jan Erik has experience on all four sides of the sourcing and vendor governance lifecycle - as a client, an industry analyst, a service provider and an advisor.

Now as a research director, principal analyst and global head of ISG Provider Lens®, he is very well positioned to assess and report on the state of the industry and make recommendations for both enterprises and service provider clients.



Provider Lens®

The ISG Provider Lens® Quadrant research series is the only service provider evaluation of its kind to combine empirical, data-driven research and market analysis with the real-world experience and observations of ISG's global advisory team. Enterprises will find a wealth of detailed data and market analysis to help guide their selection of appropriate sourcing partners. ISG advisors use the reports to validate their own market knowledge and make recommendations to ISG's enterprise clients. The research currently covers providers offering their services across multiple geographies globally.

For more information about ISG Provider Lens® research, please visit this [webpage](#).

Research™

ISG Research™ provides subscription research, advisory consulting and executive event services focused on market trends and disruptive technologies driving change in business computing. ISG Research™ delivers guidance that helps businesses accelerate growth and create more value.

ISG offers research specifically about providers to state and local governments (including counties and cities) and higher education institutions. Visit: [Public Sector](#).

For more information about ISG Research™ subscriptions, please email contact@isg-one.com, call +1.203.454.3900, or visit research.isg-one.com.

ISG (Information Services Group) (Nasdaq: III) is a leading 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 sourcing that is now at the forefront of leveraging AI to help organizations 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.

For more information, visit isg-one.com.





OCTOBER, 2026

REPORT: GENERATIVE AI SERVICES — LARGE AND MIDSIZE